

Kuzmics, Christoph; Palfrey, Thomas; Rogers, Brian W.

Working Paper

Symmetric play in repeated allocation games

Discussion Paper, No. 1551

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Kuzmics, Christoph; Palfrey, Thomas; Rogers, Brian W. (2012) : Symmetric play in repeated allocation games, Discussion Paper, No. 1551, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/97245>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Symmetric play in repeated allocation games*

Christoph Kuzmics[†], Thomas Palfrey[‡], and Brian W. Rogers[§]

July 1, 2012

Abstract

We study symmetric play in a class of repeated games when players are patient. We show that, while the use of symmetric strategy profiles essentially does not restrict the set of *feasible* payoffs, the set of *equilibrium* payoffs is an interesting proper subset of the feasible and individually rational set. We also provide a theory of how rational individuals play these games, identifying particular strategies as focal through the considerations of Pareto optimality and simplicity. We report experiments that support many aspects of this theory.

Keywords: symmetry, repeated games, focal points, experiments

*The paper has benefited from comments by participants of seminars at University of Arizona, UC San Diego, University of Cambridge, University of Iowa, University of Konstanz, Northwestern University, University of Oxford, University of Pittsburgh, Virginia Tech, University of Warwick, University of Wisconsin, and from conference presentations at the Second Transatlantic Theory Conference (2009), the 2010 Midwest Economic Theory Conference (MWIEG), the 2010 Canadian Theory Conference, and the April 2011 Midwest Political Science Association. Palfrey gratefully acknowledges the financial support from grants by the National Science Foundation (SES-0962802) and the Gordon and Betty Moore Foundation.

[†]IMW, Bielefeld University, christoph.kuzmics@uni-bielefeld.de.

[‡]Division of the Humanities and Social Sciences, California Institute of Technology, trp@hss.caltech.edu

[§]Managerial Economics and Decision Sciences, Kellogg School of Management, and NICO, Northwestern University, b-rogers@kellogg.northwestern.edu.

1 Introduction

We are interested in the economic implications of how people behave in symmetric situations with repeated interaction. We explore this issue by studying what we call n -player allocation games. An allocation game involves two issues: a coordination issue and a competition issue.

Every allocation game among n players is characterized by a length- n vector x of non-negative possible payoffs. Each player simultaneously demands an element of x . If the demands are all distinct elements of x , then payoffs are distributed according to the chosen demands. Upon any other realization, in which two or more players make the same demand, payoffs are zero to all. With exactly two players, an allocation game corresponds to the well-known battle-of-sexes game, so one interpretation of allocation games is as a generalization of the battle-of-sexes game to n players.¹

There are numerous applications for which allocation games are stylized models. Many of these have been discussed in the literature for the 2-person case, such as compatibility and standardization of technology, network externalities, product innovation, favor exchange, and so forth. For larger n , the allocation game captures important elements of role assignments in a team or organization, where efficient role assignments do not depend on idiosyncratic characteristics of the members of the team or organization. The team needs exactly one member to occupy each role in order to function profitably. Applications range from committee assignments in legislatures to allocation of routine chores in a household, to assignment of territories in a sales force, and other examples of joint production abound. The vector x captures the values or costs to the players assigned to different roles, with the simplifying assumption that joint production breaks down if any of the assignments are unfilled.

The players in an allocation game are in symmetric positions. We are thus interested in the outcomes of a repeated allocation game that obtain when play is symmetric, that is, when players use symmetric strategies. One reason to focus on symmetric play is that players may have no cues or labels outside the description of the game, as given, that could serve, via a norm or convention, to coordinate their actions on an asymmetric outcome. That is, certain environments *impose* symmetric play as a matter of technological constraint. But more broadly, symmetric equilibria play an important role in economic theory for a variety of reasons.² Thus it is important to understand the general implications of symmetric play.

Consider first a one-shot allocation game. There are of course many asymmetric (pure) equilibria of these games, but generally only one symmetric equilibrium with positive payoffs.³ But the expected payoff at the best symmetric equilibrium, which involves mixing, is low – strictly below those of the efficient asymmetric equilibria. How can players coordinate their actions, in spite of their symmetry, to achieve better payoffs? Intuitively, repetition can play an important role in such games by allowing players to converge on a long run sequence of efficient asymmetric equilibria in the absence of any exogenous coordinating device. By the same token, the inevitable inequities in allocations of the pure equilibria of the one-shot game can be smoothed out over time using, e.g., rotation schemes.⁴ In this way, the history

¹Allocation games can also be thought of as a particular n -person version of the Nash demand game, where demands are constrained to a finite set, or an assignment game in which everyone shares similar preferences over assignments.

²The notion of symmetric strategy profiles, and symmetric (Nash) equilibria, goes back to at least (Nash 1951, Theorem 2). Symmetric equilibria have been employed, sometimes more for reasons of simplicity than plausibility, in, for instance, the study of auctions, voting, and oligopolistic competition.

³There are many equilibria, both symmetric and asymmetric, that generate payoffs of zero; we are less interested in these.

⁴A familiar example is the assignment of positions in a pickup soccer or hockey game. Players must be assigned roles. Nearly everyone would rather be a forward than the other roles and nobody wants to be the goalkeeper. A common convention is to rotate players through the positions so that nobody is required to be the goalkeeper all the time and everyone spends some time at the forward position.

of the game itself can provide cues that break the symmetry and allow players to coordinate their continuation play.⁵ We study the properties of symmetric play in repeated n -player allocation games with patient players. Our goal is twofold.

First, we want to understand the implications that symmetry places on the outcomes of the repeated allocation game. What payoffs are feasible in the repeated game through the use of symmetric strategy profiles, that is, what are the *technological* constraints of symmetric play? Which of these feasible payoffs can be sustained by symmetric equilibria, that is, what are the *strategic* constraints of symmetric play?

Second, and further, we develop a theory that addresses how rational individuals play such games. There is an infinite number of subgame perfect equilibrium payoff profiles and for each one of them there may be an infinite number of equilibrium strategies that deliver those payoffs. Which payoff profile is selected? What is the particular strategy used to deliver those payoffs? We provide criteria that deliver a unique prediction, which we interpret as a *focal point* in the sense of Schelling (1960), that is, the way to play a game suggested by the structure of the game itself, in conjunction with general norms or conventions, such as Pareto optimality.⁶

The first consideration is that the focal point must be, at a minimum, a symmetric strategy profile that constitutes a subgame perfect equilibrium. The implications of this requirement are obtained through our answers to the questions above.

The next consideration is that the focal point should be Pareto optimal. We find the Pareto criterion a compelling element of a theory of focal points.⁷ Given a choice between two outcomes, Pareto optimality is precisely the condition that aligns the incentives of all agents. We find it all the more compelling in a fully symmetric environment.

To complete the theory, we leverage further the Pareto criterion by looking at the set of strategies that constitute Pareto optimal equilibria in the limit as players become perfectly patient. Among those strategies, we are able, perhaps surprisingly, to identify a particular unique strategy that is most efficient for all high discount factors.⁸ That strategy is a candidate for the focal point of the game.

As it happens, that strategy might be considered rather complex. In fact, it cannot be implemented with a finite-state automaton (although it is simple to describe). Thus we look again at the strategies that are Pareto optimal equilibria in the limit as players become perfectly patient, and find among them a simplest strategy. This strategy, which can be implemented by a very small automaton is also a candidate for the focal point of the game.

Which criteria, if any, players have in mind to determine focality is, in the end, an empirical question. Given the many possibilities, and their implications for behavior, it seems like an essential question to investigate. We report data from a battery of laboratory experiments that implement repeated allocation games with two and three players. These experiments clearly identify which of the criteria is more salient in the laboratory, and indeed there is a common focal point that emerges in all of our experimental sessions. The focal way to play also has implications for how the game is played in initial periods, that is, for the mixing probabilities in the presence of symmetry. Our experiments are designed to assess both aspects of the theory.

In the remainder of the Introduction we first (in Subsection 1.1) explain, by means of a simple example, what it means for players to use symmetric strategies in a symmetric

⁵This basic insight has its roots in models of 2-player games developed by Farrell (1987), who studied coordination via initial rounds of cheap talk in the battle of the sexes game, and Crawford and Haller (1990), in the context of pure coordination games.

⁶We are not asserting a unique theory that provides a unique prediction.

⁷There may generally be competing concerns, such as risk-dominance, but this has no bite in our setting.

⁸We characterize this strategy exactly for the case of 2-player allocation games, and discuss its extension to n -player games.

repeated game. In Subsection 1.2 we outline our results, and in Subsection 1.3 we discuss the relationship of our work to the literature.

1.1 The notion of symmetry

Consider a 2-player repeated allocation game, which is a symmetrized version of the Battle-of-the-Sexes, with stage game given by⁹

$$\begin{array}{c|cc}
 & H & L \\
 \hline
 H & 0,0 & 2,1 \\
 \hline
 L & 1,2 & 0,0 \\
 \hline
 \end{array}$$

Both players have a high action H which pays $x_1 = 2$ and a low action L which pays $x_2 = 1$. Players obtain non-zero payoffs from only the asymmetric outcomes, in which the two players take different actions, i.e. one chooses H , the other L . This stage game is played repeatedly with perfect monitoring and a common discount factor $\delta < 1$. The game is symmetric. Informally, a strategy profile is symmetric if the strategies of the two players are the same mapping from histories to actions. A first implication of symmetry is that at the empty history both players need to use the same mixed action. It cannot be, for instance, that in the first stage one player chooses H (with probability one) and the other player chooses L (with probability one), i.e., the efficient equilibria are precluded by symmetry. In the second stage there are three possible histories: both players played H , both players played L , and the players took distinct actions in the first stage. In the first two cases, the players must use the same mixed action as each other in the second stage. In the terminology of Bhaskar (2000), which we shall also employ, symmetry has not been broken. Notice though, that this mixed action, used by both players, can be different from the mixed action used in the first stage, and can depend as well on whether the history is (H, H) or (L, L) . In the remaining case, symmetry is broken: it is common knowledge between the players that they have taken different actions, and this fact distinguishes them. At this history, and forever after, the two players can use different distributions over actions. Symmetry implies, however, that whatever mixed action is taken in the second stage by the player who in the first stage played H , would have also been taken by his opponent, had the first stage realizations been exchanged. Of course, it could happen that symmetry is broken only later, say in stage T . Then it must be the case that in every stage preceding T both players employ the same distribution over actions. Finally, at stage T it must have been that, for the first time in the game, one player's action realized as L , while the other's realized as H . Because symmetry means only that the players map histories to future actions in the same way, once symmetry is broken, joint continuation play is unrestricted at subsequent stages.

Naturally, if two players employ the same strategy in a symmetric game they must expect the same payoff. However, on the path of play they may very well obtain different payoffs. For instance, they could randomize $\frac{1}{2}H + \frac{1}{2}L$ until symmetry is broken and then, once an (H, L) has realized, continue to play those same actions (H, L) forever after. Thus, while they use symmetric strategies and expect the same payoff ex-ante, they also expect a difference in their payoffs. In other words, conditional on being the H -player, i.e. conditional on being the player who plays H at the stage when symmetry is broken, a player expects a continuation payoff of 2, which is different than the payoff of 1 that the L -player, the player who plays L when symmetry is broken, expects. In this regard we shall refer to the *payoff of the H -player*, with which we mean the expected discounted payoff to a player conditional on this player being the H -player as just described. Similarly we define the *payoff of the L -player*.

⁹This example also motivates Bhaskar (2000), which in turn motivated us. We discuss Bhaskar's (2000) work and its relation to our paper in Subsection 1.3.

Symmetry ensures these are well-defined, as payoffs do not depend on the identity of the player who so happens to become the H -player along a path of play.

We shall say that a symmetric strategy profile is *ex-post payoff-symmetric* if the expected payoff (discounted back to time 0) conditional on being the H -player is the same as the payoff conditional on being the L -player.¹⁰ More generally, ex-post payoff-symmetry requires that every player expects the same payoff independent of which action he takes when symmetries are broken, which is to say that each player expects no difference between his payoff and the payoff of any other player. Observe that symmetric strategy profiles are not generally ex-post payoff-symmetric.

1.2 The Results

Our first result answers the following question. Could it be that restricting to symmetric strategy profiles in itself implies that play can never yield highly asymmetric payoffs? That is, does the restriction to symmetric strategies directly imply restrictions about the symmetry of payoffs. If so, one might argue on epistemic grounds that symmetric strategies are too strong of a theoretical restriction. However, we show this is not the case; the answer to this question is emphatically no. Proposition 1 establishes that symmetric strategy profiles can generate essentially all feasible payoff profiles for patient players. For example, in the 2-player battle-of-the-sexes game from Subsection 1.1, this means that a symmetric strategy profile can induce essentially any long run payoff pair that an arbitrary strategy profile can induce, where one payoff is interpreted as the payoff to the H -player player and one as the payoff to the L -player.¹¹

We then turn to the properties of equilibrium payoffs obtained under symmetry. Specifically, we ask the following question: How inefficient are symmetric equilibria that are not ex-post payoff-symmetric? In particular if the discount factor is close to one, could it be that ex-post payoff-asymmetric equilibria are nearly efficient (i.e. close to optimal)? There are reasons to believe that they might be. Suppose players are promised an asymmetric payoff profile at the point when symmetries are broken. The payoff asymmetry implies that the players will compete to obtain the higher payoffs by mixing with large probability on those actions. This will result in substantial delay in obtaining coordination. As the discount factor tends to one, there are two forces at work. On one hand the expected delay becomes longer, but on the other hand players are more patient and so delay of a given length becomes less costly. Which of these effects dominates?

We obtain two results that directly address these questions. First, for any symmetric equilibrium strategy profile that is not ex-post payoff-symmetric there is a distinct loss of efficiency relative to the optimal symmetric equilibrium, even as the discount factor approaches one. This is established in Proposition 2. Therefore, not all feasible and individually rational payoffs can be supported by a symmetric equilibrium. There is a meaningful restriction, in this sense, on the set of payoff profiles. Some degree of asymmetry in payoff profiles can be sustained in a symmetric equilibrium, but there is a general tension between asymmetries in ex-post payoffs and the corresponding level of efficiency that can be sustained in equilibrium.

Using a simple version of the 2-player allocation game, Propositions 3 and 4 then characterize the set of payoff profiles sustainable in a symmetric equilibrium of the repeated game and prove that it is a relatively small subset of the set of feasible and individually rational payoff profiles, with an intuitive structure. It is relatively small in the sense that it does not include any payoff-profiles that are both nearly efficient and highly asymmetric. However, it is large enough to include payoff profiles that are both moderately asymmetric and moderately inefficient.

¹⁰Bhaskar (2000) calls such a strategy profile egalitarian.

¹¹This is precisely true only in the limit in which the discount factor tends to one.

Finally, in a corollary to Proposition 2, we show that there is a *unique most efficient payoff profile* sustainable by a symmetric equilibrium, even when the discount factor tends to one. That is, Pareto efficiency selects a unique payoff profile, even with perfectly patient players.¹²

However, while Pareto optimality thus pins down a particular payoff profile, there remains an infinite number of symmetric equilibria that generate the efficient payoffs. We thus turn to the final question as to which of the efficient symmetric equilibria is focal for patient players. We look for a symmetric equilibrium that is not only ex-post payoff-symmetric in the limit, but also most symmetric among all strategies that are ex-post payoff-symmetric in the limit, for all high discount factors. Perhaps surprisingly, there does exist a strategy profile with this feature.¹³ It is such that continuation play at every history at which symmetry is broken is given by the Thue-Morse sequence.¹⁴ This result is established in Proposition 6. Proposition 5 characterizes the Thue-Morse sequence in economic terms, as the limit sequence that is obtained for patient players, when play is such that at every stage, the player with the currently lower total discounted payoff receives the highest stage game payoff.

In order for a particular strategy profile to be focal, it must be sufficiently simple that, not only can agents identify it, but (at a minimum) they must have faith that their opponents can do so as well. We want our theory to be descriptive and, as such, simplicity forms the final element of our approach. While the strategy profile that gives rise to the Thue-Morse sequence once symmetries are broken can be represented by a finite Turing-machine, it cannot be represented by a finite automaton. If, in the minds of players, simplicity of a strategy is more appropriately captured by the state-complexity of the smallest automaton that represents this strategy, one must abandon the Thue-Morse continuation play in favor of a simpler strategy. As the discount factor tends to one, there is a simplest strategy in the automaton sense that induces efficient payoffs in equilibrium. This strategy adopts a rotation scheme among the possible payoffs of x once symmetry is broken. These results are formally stated as Propositions 7 and 8.

The theory suggests that symmetric payoffs can be justified purely through an efficiency criterion, without any inherent preference for “fairness”. Delivering symmetric payoffs, though, requires a greater level of complexity than delivering asymmetric payoffs, in that players have to manage to equally share the payoffs following the breaking of symmetry, by use of a scheme that delivers different stage game payoffs over time. Do players manage to do this? Our experiments demonstrate that the answer is, overwhelmingly, “yes”. More specifically, nearly symmetric payoffs are almost always delivered, and they are delivered via rotation schemes, even in the more difficult 3-person games, where multiple efficient rotation schemes are possible. We interpret this as strong evidence that Pareto efficiency and, secondarily, simplicity, are elements of focal strategies in the lab. Further, this result has implications for how players should mix before symmetry is broken. In particular, anticipating ex post symmetric payoffs, players should mix close to uniformly, thereby breaking symmetry as fast as possible. This implication about mixed strategies prior to coordination is also supported by the data.

1.3 Related literature

Bhaskar (2000) studies 2-player allocation games, where the mis-coordination payoffs are not necessarily identical. Using this stage game he studies the finitely repeated game with and

¹²This payoff profile is ex-post payoff-symmetric. Bhaskar (2000) establishes a similar result for 2-player repeated allocation games that are either finitely repeated or have discount factor strictly less than one.

¹³For any fixed $\delta < 1$, there exist exactly ex-post symmetric strategies (see Lemma 1). But these strategies are highly sensitive to small changes in δ .

¹⁴See (Thue (1906), Thue (1912), and Morse (1921). Allouche and Shallit (1999) characterize some interesting mathematical properties of this sequence.

without discounting and the infinitely repeated game with discounting. In his Propositions 1 and 2 he shows that (under some specific conditions on the payoff matrix and provided the discount factor is large enough) “in any optimal symmetric equilibrium of [any of these repeated games] the players adopt [an] egalitarian convention” in the event that symmetry is broken. In other words any symmetric equilibrium that provides the highest possible ex-ante payoff to the two players (among all symmetric equilibria) must be ex-post payoff-symmetric.

The basic insight behind this result is an important one: Any symmetric strategy profile that, upon the breaking of symmetry, promises non-egalitarian payoffs (e.g. higher payoffs to the H - than to the L -player) induces inefficient initial randomization as incentives are skewed in favor of the action that provides the higher ex-post payoff. Optimally the initial randomization should be uniform.

This result for 2-player games builds on two earlier related contributions. The first is Farrell (1987), who studies the dynamic game in which players engage in multiple rounds of strategic cheap-talk prior to playing a one-shot symmetric version of the battle-of-the-sexes, i.e. essentially the game of our Subsection 1.1. In every round of the cheap-talk phase players can send one of two messages (messages are labeled as the eventual actions in the final stage game). Farrell (1987) shows that no matter how many rounds of cheap-talk are available to the two players, they will not coordinate with probability 1 in the final game. The mechanism that drives this result is similar to that which drives Bhaskar’s (2000). In the optimal equilibrium, the person who is the first in the cheap-talk phase to be alone when sending the high-payoff action as her message, is the one that will play the high-payoff action in the battle-of-the-sexes game. This induces both players to randomize between the two messages with a much larger weight on the high-payoff action. The other related contribution is Crawford and Haller (1990), who, as far as we know, were the first to properly study the implications of symmetry in repeated games. Both Bhaskar (2000) and our paper owe a methodological debt to Crawford and Haller (1990). However both Bhaskar (2000) and this paper differ from Crawford and Haller (1990) in that Crawford and Haller (1990) study repeated pure coordination games, without the competition issue inherent in the allocation game. Blume (2000) further develops Crawford and Haller’s (1990) approach to study aspects of language.

Blume and Gneezy (2000) and Blume and Gneezy (2010) perform lab experiments with symmetric coordination games in the spirit of Crawford and Haller (1990). Blume and Gneezy (2000) demonstrate that for a simple twice repeated game of pure coordination (with 3 actions) many subjects play according to the optimal symmetric strategy profile. In a more complex twice repeated game also of pure coordination (with 9 actions) subjects mostly fail to play the optimal symmetric strategy profile. In Blume and Gneezy (2010) further experiments are undertaken that demonstrate that there are two reasons for the failure of optimal play, one is that some subjects cannot identify the optimal way to play and the other is that some subjects who are able to identify optimal play do not have sufficient faith that their opponents are able to do so.

The notion of rotation schemes in repeated interactions dates back to Luce and Raiffa (1957). Lau and Mui (2008) and Lau and Mui (2012) study a particular equilibrium in 2-player repeated games involving “turn taking” as a way to implement nearly symmetric payoffs, which is a specific instance of our rotation schemes. Prisbey (1992) is the first paper to observe rotation in experiments, again in the context of 2-player games. Cason, Lau, and Mui (2012) observe rotation schemes in 2-player assignment games, demonstrating that subjects are able to teach such behavior to future opponents.

The paper proceeds with a formalization of the model in Section 2. Feasible payoffs are described in Section 3. The main theoretical result, which describes the equilibrium tension between asymmetric and efficient payoffs, is presented in Section 4. There, we also show with an example that very asymmetric payoffs are necessarily very far from efficient. Section 5 identifies the focal strategy that uses only Pareto optimality for high discount factors.

Section 6 identifies rotation schemes as being the simplest nearly efficient strategies. We report the details of the experimental design and the results of the experiments in Section 7. Some of the longer and more technical proofs, and an example with a different symmetry structure, are relegated to the Appendix.

2 Model

2.1 The stage game

While many of our results can be generalized to other settings, we focus attention to a particular class of games.

Definition 1. A symmetric normal form stage game $\Gamma = (I, A, u)$ is an n -player allocation game if $A = \{1, \dots, n\}$, i.e. $|A| = |I| = n$, $u(a) = (0, \dots, 0)$ for all action profiles $a \in A^n$ with the property that there are two players $i \neq j$ such that $a_i = a_j$, and there exists an $x \in \mathbb{R}^n$ such that (i) $x_1 > \dots > x_n \geq 0$, and (ii) $u_i(a) = x_{a_i}$ for all $i \in I$ if a is a permutation of $(1, 2, \dots, n)$.

That is, every one of n players requests any one of n amounts of money.¹⁵ If demands are consistent, such that every player requests a different amount, then players are paid their respective demands. Otherwise players are paid zero. Note that all players in an allocation game are symmetric, while there are no symmetric strategies. Note that for $n = 2$ this reduces to the battle-of-the-sexes.

2.2 The repeated game

The game is played repeatedly at discrete points in time $t = 0, 1, 2, \dots$. Players discount future payoffs with a common discount factor $\delta < 1$. In each period players observe an element of $Y = \{c, n\}$ (as well as their own realized action). If $a^t \in A^n$ is played at stage t then players observe c if a^t is a permutation of $(1, 2, \dots, n)$. Otherwise they observe n . Thus c is the “event” that players achieved coordination (a non-zero payoff vector) and n is the event that they did not achieve coordination and thus obtained 0 payoffs all. The payoff matrix is assumed to be common knowledge. Players, thus, know what payoff they received at each stage. Note that for 2-player allocation games, this information structure is equivalent to perfect monitoring. Otherwise, monitoring is less than perfect.

We can thus describe public and private histories for the repeated game. The set of public histories is given by $\mathcal{H} = \bigcup_{t=0}^{\infty} Y^t$ with $Y^0 = \emptyset$. Player i 's set of private histories is given by $\mathcal{H}_i = \bigcup_{t=0}^{\infty} (A \times Y)^t$. Given the symmetry we have $\mathcal{H}_i = \mathcal{H}_j = \mathcal{H}^*$ for all $i, j \in I$. A pure strategy is a mapping $\sigma_i : \mathcal{H}^* \rightarrow A$. A behavioral strategy is a mapping $\sigma_i : \mathcal{H}^* \rightarrow \Delta(A)$. For convenience we shall also define the set of all play paths by $\mathcal{P} = (A^n)^\infty$.

Finally we need to specify payoffs. For a pure strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ payoffs are given by $u_i(\sigma) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t(\sigma))$, where $a^t(\sigma)$ is the action profile induced by strategy profile σ in period t . For mixed strategy profiles we extend u_i by taking expectations.

¹⁵In these games we restrict attention to payoff vectors that have distinct elements. This is not a crucial assumption for our results. However, it implies that, while all players are symmetric, there are no symmetric stage game actions. This avoids expository complications. See Crawford and Haller (1990), where symmetry among actions is a central element of the model and analysis.

2.3 The solution concept

The key departure we make in this paper from the usual analysis of repeated games is that we restrict attention to symmetric strategy profiles.¹⁶ The requirement of symmetric play among symmetric players is perhaps best explained as follows. The most compelling justification for Nash equilibrium is that if a public recommendation is made as to how to play a game, this recommendation must satisfy the conditions of a Nash equilibrium in order to potentially be realized. Now suppose that in a symmetric game, in which players have no commonly known names (or commonly known distinguishable characteristics), a recommendation is made that one player should play one strategy and another player should choose another strategy.¹⁷ Then, as players cannot be called by name, they would have to figure out for themselves who plays which component of the recommendation. If they then cannot communicate it will not be possible for them to follow such a recommendation without the positive probability event of both choosing the same strategy. Thus recommendations in the presence of player symmetry must describe a single strategy that all symmetric players can use, in order to have any hope that play will follow the prediction. Of course, the recommendation must also be a (sequential or subgame perfect) Nash equilibrium. Otherwise each player will have an incentive to deviate (at some history). Thus, we shall look at symmetric subgame perfect Nash equilibria of the repeated game.

Formally, given the notation above, a behavioral strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is symmetric if, simply $\sigma_i = \sigma_j$ for all $i, j \in I$. Note that this does not imply that two players necessarily behave in the same way after a particular given public history, even if symmetries have not yet been broken, as they may well have different private histories. If two players, up to a particular history, ended up playing the same action in every stage, say both played action a in stage 1, both action b in stage 2, and so on, then they must indeed play the same mixed action in the stage following that history. If, on the other hand, they have not played in identical fashion up to a given history¹⁸, then their play after that can differ arbitrarily. Note, however, that there is an additional restriction imposed by symmetry. If player 1 played, say, actions a and then b , while player 2 played b and then a , for a given 2-period public history, then player 1 has to behave in the stage after that just as player 2 would have done had she played a and then b and player 1 had played b and then a , and given that same public history.

2.4 Notions of payoff symmetry

Consider an n -player allocation game $\Gamma = (I, A, u)$. Consider a behavioral strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$. This strategy profile σ induces a probability distribution over the set of all play paths \mathcal{P} . Recall that $u_i(\sigma)$ denotes player i 's expected (discounted) payoff from the repeated game strategy profile σ .

Definition 2. A strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is *ex-ante payoff-symmetric* if $u_i(\sigma) = u_j(\sigma)$ for all $i, j \in I$.

¹⁶More generally, we would be interested in strategy profiles that respect the, possibly partial, symmetry structure of the game. Such strategy profiles have been termed “attainable” by Crawford and Haller (1990) in their study of repeated pure coordination games and not simply “symmetric”, to emphasize that the game does not have to be fully symmetric for symmetry restrictions to bite. Restrictions of symmetries in the game, following Crawford and Haller (1990), have been further studied by Alos-Ferrer and Kuzmics (2008). See also Harsanyi and Selten (1988) and Casajus(2000, 2001) for related notions.

¹⁷The requirement that the players have no commonly known names or characteristics is important. If they do – for instance it is commonly known that one of the players is a man, while the other is a woman – one can announce that the woman do one thing, while the man do another.

¹⁸They, thus, must have used a mixed behavioral strategy with different action realizations.

If σ is a symmetric strategy profile then it is ex-ante payoff-symmetric. Not every ex-ante payoff-symmetric strategy profile is also symmetric.

For a given play-path $p \in \mathcal{P}$ denote $T(p) = \min_{t \geq 0} \{t \mid a_i^t \neq a_j^t \text{ for all } i, j \in I\}$, the time period in which symmetry is broken along this path. Consider player $i \in I$. Consider all play-paths in which player i plays action $a \in A$ when symmetries are broken. This set is given by $\mathcal{P}_i(a) = \{p \in \mathcal{P} \mid a_i^{T(p)} = a\}$.

Define player i 's (discounted repeated game) payoff conditional on player i being the a -player, that is, conditional on the event $p \in \mathcal{P}_i(a)$, by $w_i(a, \sigma) = \mathbb{E}[u_i(\sigma) \mid p \in \mathcal{P}_i(a)]$.

Note that if $\sigma = (\sigma_1, \dots, \sigma_n)$ is symmetric, then, necessarily, $w_i(a, \sigma) = w_j(a, \sigma)$ for all $i, j \in I$ and for all $a \in A$. In this case, we can drop the subscript and denote by $w(a, \sigma)$ the expected payoff of the a -player, i.e. the player that plays a at the time when symmetries are broken.

Definition 3. A symmetric strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is ex-post payoff-symmetric if $w(a, \sigma) = w(b, \sigma)$ for all $a, b \in A$.

It is not true that every symmetric strategy profile is ex-post payoff-symmetric (as we shall show in the next Section). There are also stronger notions of ex-post symmetry. Note that the event that determines our notion of ex-ante and ex-post is the event that symmetry is broken. Now there may well still be some randomness, induced by σ , after the event that symmetry is broken. One might want to strengthen the notion of ex-post symmetry, call it *strong ex-post payoff-symmetry*, to demand that the payoffs to all players (differentiated by the action they took when symmetries were broken) are equal to each other for every realization of the remaining path of play. Of course, strong ex-post symmetry implies ex-post symmetry. If σ prescribes a deterministic play-path after the event when symmetry is broken then ex-post symmetry implies strong ex-post symmetry. If, however, there are still multiple play-paths probable after the event when symmetry is broken then such a strategy profile may well be ex-post payoff-symmetric, yet not strongly so. For instance¹⁹, suppose the game is a 2-player allocation game and the continuation play (induced by σ) after symmetry is broken is such that the two players observe a public randomization device (often employed as a short-cut to prove folk-theorems in repeated games) that is a fair coin toss. If the coin comes up heads the strategy 1-player gets to play the action associated with the high payoff, while the strategy 2-player has to play the action with the low payoff. If the coin comes up tails, they do the opposite. Such a strategy profile is ex-post payoff-symmetric as at the time of symmetry being broken both players do not yet know the outcome of the coin-toss, but is not strongly ex-post payoff-symmetric, as eventually they get different payoffs along the realized path of play.

It turns out that for our setting of repeated allocation games, without allowing public randomization devices, if a symmetric strategy profile is an efficient equilibrium it must be strongly ex-post payoff-symmetric, while for general symmetric n -player games one can only show that if a symmetric strategy profile is an efficient equilibrium it must be ex-post payoff-symmetric. That such strategy profiles do not generally have to be strongly ex-post symmetric, even without public randomization devices, is shown by example in Appendix D.

This distinction between ex-post payoff-symmetry and strong ex-post payoff-symmetry also becomes important in the interpretation of discounting. If discounting is such that players are actually facing a shrinking pie then the two notions are more closely related than if discounting is such that there is, at every stage, a certain probability that play stops. While, in a certain sense, the two interpretations of discounting are strategically equivalent,²⁰ a strategy profile may well be strongly ex-post payoff-symmetric under the

¹⁹We thank Jeff Ely for pointing this out.

²⁰As discussed further in Section 7, we implement both conditions in the laboratory and find no difference in play.

shrinking pie interpretation but not strongly ex-post payoff-symmetric under the random stopping interpretation. In the latter case there are, by nature, still many more play-paths possible after symmetry is broken than in the former case.

3 Feasible Payoffs

Consider a repeated n -player allocation game. Consider the following strategy, represented in automaton form, $(\mathcal{W}, w^0, f, \tau)$.²¹ Let $w^0 = R \in \mathcal{W}$ be the initial state. Let $f(R)$ be the uniform distribution over all actions in A . Let $\tau(R, a, y) = R$ if $y = n$ (symmetry was not broken), and $\tau(R, a, y) = S(a)$ for some $S(a) \in \mathcal{W}$ with $S(a) \neq S(b)$ for all $a, b \in A$ if $y = c$ for the first time (symmetry was broken). This automaton represents a strategy in which all players initially randomize uniformly over all actions in every period until coordination is achieved (i.e. all players use a different action) and, hence, symmetry is broken. After that, if all players use this automaton, all players' automata will now be in different states. Thus, from that point on they can, in principle, play any (possibly asymmetric) strategy profile of the repeated game. Let us assume that the continuation played does not depend on the particular history at which symmetry is broken. Thus, let $v(a)$ denote the continuation payoff for the player who played action $a \in A$ when symmetry was broken. We will show that this class of strategies is sufficient to generate essentially all feasible payoffs of the repeated game without symmetries.

Let $v \in \mathbb{R}^n$ denote the vector of these continuation payoffs. Given the uniform distribution over actions the probability of coordination in any given stage is $q = \frac{n!}{n^n} > 0$. Thus, players will eventually coordinate, symmetry will be broken, and players obtain the payoff-vector v from the continuation play in the now unrestricted repeated game. Let $w(a)$ denote the ex-ante expected payoff to the player who eventually plays action a when symmetry is broken. Then

$$w(a) = [(1 - \delta)x_a + \delta v(a)] q \sum_{t=0}^{\infty} \delta^t (1 - q)^t = [(1 - \delta)x_a + \delta v(a)] \frac{q}{1 - \delta(1 - q)}.$$

Let the set of payoff-vectors that are feasible through the use of all (including asymmetric) strategy profiles of the repeated game be denoted by $\mathcal{F}_\Gamma \subset \mathbb{R}^n$. Note that \mathcal{F}_Γ , while in principle dependant on the discount factor δ , is actually constant for all $\delta \geq \bar{\delta}$ for some $\bar{\delta} < 1$. This follows from a result in Sorin (1986), also stated as Lemma 1 in Fudenberg and Maskin (1991) and as Lemma 3.7.1 in Mailath and Samuelson (2007). In the 2-player allocation game, in fact, we have that $\mathcal{F}_\Gamma(\delta) = \mathcal{F}_\Gamma$ for all $\delta \geq \frac{1}{2}$. Let $\mathcal{F}_\Gamma^s(\delta) \in \mathbb{R}^n$ denote the set of feasible payoff-vectors under symmetric strategy profiles. A typical element $w \in \mathcal{F}_\Gamma^s(\delta)$ is, thus, a vector of payoffs, in which each coordinate corresponds to an action $a \in A$ and represents the expected discounted payoff, $w(a)$, for some underlying symmetric strategy profile, to a player conditional on this player being the a -player (the player who plays action a when symmetry is broken). Then

$$\mathcal{F}_\Gamma^s(\delta) \subset \left\{ w \in \mathbb{R}^n \mid w = [(1 - \delta)x + \delta v] \frac{q}{1 - \delta(1 - q)} \text{ for some } v \in \mathcal{F}_\Gamma \right\}.$$

It is thus obvious that, as δ tends to 1, i.e. as players become increasingly patient, the set of feasible payoffs under symmetric strategies coincides with the set of feasible payoffs under all strategies.

²¹The set \mathcal{W} is the set of states for the automaton, w^0 is the initial state, $f : \mathcal{W} \rightarrow A$ is the (possibly random) action chosen as a function of states (which could also depend on observables from the game, but this is not needed for our purposes), and finally $\tau : \mathcal{W} \times A \times Y \rightarrow \mathcal{W}$ is the state transition function.

Proposition 1. *Let Γ be an n -player allocation game. As the discount factor tends to one, the set of payoff profiles that are feasible under symmetric strategies tend to the set of feasible (unrestricted) payoff profiles. That is*

$$\lim_{\delta \rightarrow 1} \mathcal{F}_\Gamma^s(\delta) = \mathcal{F}_\Gamma.$$

Note that Proposition 1 immediately generalizes to all symmetric n -player games as long as it is possible to break all symmetries in the game with some symmetric strategy profile with probability 1 in finite time.²²

4 Equilibrium Payoffs

While symmetry, thus, hardly poses a restriction on the set of feasible payoffs, this section demonstrates that symmetry does impose interesting restrictions on the set of equilibrium payoffs for all discount factors, even as δ tends to 1.

Consider a repeated n -player allocation game with coordinated payoff vector x . Recall that x_a is the stage payoff a player receives when playing a and all other players play in such a way that no two players choose the same action. We can translate this payoff vector of the stage game to payoff profiles in the repeated game as follows. Let Π be the space of action permutations with typical element $\pi : A \rightarrow A$. Let w^π be a payoff profile such that $w^\pi(a) = x_{\pi(a)}$, i.e. the repeated game payoff to the a -player is given by the stage game payoff a player gets from playing action $\pi(a)$ (in the action profile in which no two players play the same two actions).

Note that for all π we have $w^\pi \in \mathcal{F}_\Gamma$, and, thus, by Proposition 1, also in $\lim_{\delta \rightarrow 1} \mathcal{F}_\Gamma^s(\delta)$. This is to say that payoff profile x (and all its permutations) are feasible under a symmetric strategy profile, at least in the limit in which the discount factor tends to 1. Furthermore, there are many payoff profiles y that are close to some w^π that are also feasible under symmetric strategy profiles for δ sufficiently close to 1.

Denote by $\mathcal{Z}_\Gamma = \text{conv}\{w^\pi\}_{\pi \in \Pi}$ the Pareto frontier of \mathcal{F}_Γ , that is, the convex hull of the permutations of x .

For a given payoff z and a small positive number $\epsilon > 0$ let $\mathcal{U}_\epsilon^z = \{y \in \mathcal{F}_\Gamma \mid \|y - z\|_\infty \leq \epsilon\}$, where $\|\cdot\|_\infty$ is the infinity norm (the difference between y and z is the maximal difference between their coordinates). Note that for any $z \in \mathcal{Z}$ and any $\epsilon > 0$ there is a $\bar{\delta} < 1$ such that for all $\delta \geq \bar{\delta}$ we have that $\mathcal{U}_\epsilon^z \cap \mathcal{F}_\Gamma^s(\delta) \neq \emptyset$.

We say that a strategy is Markov if it has two properties: (i) for every history at which symmetry is not broken, the mixed action is identical, and (ii) the continuation play upon the breaking of symmetry depends only on the realized action at the stage in which symmetry is broken. Let $\mathcal{E}_\Gamma^s(\delta)$ denote the set of payoff profiles (as a vector to the respective a -players) that are sustainable under symmetric *Markov equilibria* of the repeated game.²³

Proposition 2. *Consider any repeated n -player allocation game Γ and any Pareto efficient payoff profile $z \in \mathcal{Z}_\Gamma$ that is not completely symmetric, i.e. such that there exist a and b for which $z_a \neq z_b$. There is an $\epsilon > 0$ such that for all $\delta \in (0, 1)$ we have $\mathcal{U}_\epsilon^z \cap \mathcal{E}_\Gamma^s(\delta) = \emptyset$.*

²²There are pathological games in which symmetries can never be broken. For instance, consider the repeated 2-player allocation game with random stopping, without any feedback, and payoffs being paid only after the game ends. In this game, there is never any information on which players can coordinate.

²³The result applies only to Markov equilibria. We conjecture that the same conclusion holds for all symmetric equilibria. Proving so may rely, in part, on characterizing the optimal way for n players to break symmetries in the repeated game. The proof for Markov equilibria is short and elegant as it uses only a few key aspects of the structure of symmetric equilibria. Extending the proof to all symmetric equilibria seems to necessitate a more complicated proof accommodating detailed aspects of the equilibrium structure. We hope to prove this extension.

In words Proposition 2 states that, for any asymmetric payoff profile that is close to the Pareto frontier, there is no symmetric Markov equilibrium that supports those payoffs for any discount factor, even as δ tends to one. Its proof is given in Appendix A. We provide here an intuition for the result. Consider a symmetric Markov strategy profile that is not ex-post payoff-symmetric, but for which the payoffs are close to the Pareto frontier. The proposition focuses on the following unilateral deviation. Consider playing, at every round before symmetry is broken, an action which promises the highest continuation payoff (and then playing as prescribed in the continuation game). How much will this cost in terms of slowing down the process of breaking symmetry? Since payoffs are close to efficient, it must be that players are breaking symmetry fast, relative to the discount factor. Using this fact, the proof shows that the delay cost of the deviation can be made as small as needed through the choice of ϵ . Now, since the strategy profile is not ex-post payoff-symmetric, the deviation also promises the player a strictly higher continuation payoff once symmetry is broken, thus making it profitable.

We remark that ϵ cannot be chosen uniformly with respect to payoff profiles $z \in \mathcal{Z}_\Gamma$ along the Pareto frontier. Let us define $z^* \in \mathcal{Z}_\Gamma$ to be the (unique) payoff-symmetric point on the Pareto frontier, i.e., $z_a^* = 1/n \sum_{a \in A} x_a$ for all $a \in A$. Roughly speaking, the more symmetric is z , the smaller ϵ must be chosen, with the particular property that as z approaches z^* , the associated choices of ϵ must converge to zero.

This observation has two important implications. First, there is a unique efficient payoff profile that can be sustained by a symmetric equilibrium as δ tends to one, and that profile is the one that is exactly ex-post symmetric. We have the following.

Corollary 1. $\mathcal{Z}_\Gamma \cap \lim_{\delta \rightarrow 1} \mathcal{E}_\Gamma^s(\delta) = z^*$.

In words, as δ tends to one, the only efficient payoff profile sustained by a symmetric equilibrium is that which is ex-post payoff-symmetric.

In light of Proposition 2, which implies that the intersection cannot be bigger than z^* , this can be proven by construction. Consider a strategy profile that promises continuation payoffs z^* from every history at which symmetry is broken. It is easily verified that uniform mixing over A at all preceding histories constitutes an equilibrium. So as $\delta \rightarrow 1$, this equilibrium delivers payoffs z^* .

The second implication is that payoff profiles that are more asymmetric necessitate more efficiency loss to implement. This reduction in equilibrium efficiency will tend to be severe as z approaches the least symmetric points on the Pareto frontier, given by w^π . While Proposition 2 does not precisely characterize this tradeoff for arbitrary allocation games, the next section demonstrates the magnitude of this effect by means of the simplest possible example.

We conjecture that symmetric repeated games outside the class of n -player allocation games typically satisfy the property that efficient payoffs implies ex-post payoff-symmetry (without the strong qualification). Appendix D provides an example of a symmetric repeated game outside our class in which an efficient symmetric equilibrium, while being ex-post payoff-symmetric, is not strongly so.

4.1 A 2-player example

Proposition 2 demonstrates that, for any n -player allocation game, there is a meaningful qualitative restriction on behavior imposed by symmetric equilibrium. It does not, however, provide a big quantitative restriction on the set of symmetric equilibrium payoff profiles. To accomplish this we investigate in this subsection a particular 2-player allocation game.

We show two things. First, the restriction imposed by symmetric equilibrium on payoff profiles is substantial. No player, in the repeated game, can obtain a payoff that exceeds

$\frac{3}{4}$ of the total available stage game payoff. Second, the restriction imposed by symmetric equilibrium is not so severe as to eliminate a Lebesgue measure 1 set of payoff profiles.

Consider the 2-player allocation game with $x_1 = 1$ and $x_2 = 0$. I.e. the stage game is given by

	H	L
H	0,0	1,0
L	0,1	0,0

where we label the actions $A = \{H, L\}$. Note that H (for “high”) weakly dominates L (for “low”) within the stage game.²⁴ As with all 2-player allocation games, observing c or n , i.e. whether or not coordination has been achieved, along with one’s own action, is sufficient for players to know exactly what has been played. Thus, the repeated game is one of perfect monitoring.

Denote, as before, by \mathcal{F} the set of payoff vectors that are feasible when considering all strategy profiles of the repeated game, and denote by $\mathcal{F}^s(\delta)$ those payoff vectors that are feasible under the restriction that strategy profiles be symmetric.²⁵ The set \mathcal{F} does not depend on δ , as long as $\delta \geq \frac{1}{2}$, and is equal to the triangle given by the convex hull of payoff pairs $(0, 0)$, $(0, 1)$, and $(1, 0)$. The set $\mathcal{F}^s(\delta)$ is a sub-triangle, which, by Proposition 1, tends to \mathcal{F} as δ tends to 1.

It is well-known and, in fact, immediate for the game at hand, that the set of equilibrium payoff pairs, without the symmetry qualification, is equal to \mathcal{F} . Denote by $\mathcal{E}^s(\delta)$ the set of symmetric equilibrium payoff pairs (here without the Markov requirement). An element of $\mathcal{E}^s(\delta)$ is a pair of payoffs, the first is the expected discounted payoff to the H -player, the second the expected discounted payoff to the L -player. We shall show that $\mathcal{E}^s(\delta)$ includes many payoff-pairs that are not ex-post payoff-symmetric, but yet is not nearly as large as $\mathcal{F}^s(\delta)$.

To obtain an interesting upper bound on the set of symmetric equilibrium payoff pairs in this repeated game, we shall appeal to a fixed point argument of an appropriate function in the set of potential symmetric equilibrium payoff pairs. Consider time 0 or any time period in which symmetry has not yet been broken. The potential outcomes of play in this stage are the pure strategy combinations HH , HL , LH , and LL .

Suppose HL occurs. Then the H -player obtains a payoff of 1, the L -player one of 0, and symmetry is now broken. Thus, the continuation from here on can be any (possibly asymmetric) equilibrium of the repeated game. Thus, the continuation payoffs can be any element in \mathcal{F} , the set of feasible payoff-pairs.

We can, in fact, combine the (relatively negligible) one-period payoff and the later discounted per-period continuation-payoff in \mathcal{F} , by assigning the two players an appropriate payoff-pair in \mathcal{F} right at this time 0 (upon observing outcome HL). That is, upon observing outcome HL , continuation payoffs can be (almost) any pair $(w^H, w^L) \in \mathcal{F}$, where $w^H, w^L \geq 0$ and $w^H + w^L \leq 1$, denoting by w^a the expected discounted payoff to the a -player.²⁶

After outcomes HH and LL both payers receive a stage payoff of 0 and symmetry is not yet broken. Thus, the continuation payoffs given HH or LL can be any symmetric equilibrium payoff-pair. Note, however, that play does not necessarily have to continue in the same way after HH and LL .

Thus, after HH a continuation might be any $(w_{HH}^H, w_{HH}^L) \in \mathcal{E}^s(\delta)$, which, as we shall see, is a more severe restriction than just $w_{HH}^H, w_{HH}^L \geq 0$ and $w_{HH}^H + w_{HH}^L \leq 1$. After

²⁴This dominance relation between the two actions, however, is not important in the analysis and does not drive the results. The analysis is similar if the zeros are replaced with a small positive number.

²⁵As we consider a particular game in this subsection, we drop the dependence on Γ in the notation.

²⁶This is precisely true in the limit as δ tends to one, which is the case we consider. For $\delta < 1$, it is not quite possible to, e.g., assign the H -player a continuation payoff of zero.

LL continuation payoffs might also be any $(w_{LL}^H, w_{LL}^L) \in \mathcal{E}^s(\delta)$, possibly different from (w_{HH}^H, w_{HH}^L) .

Summarizing, at time 0 players expect the following continuation profiles: $(w^H, w^L) \in \mathcal{F}$ to the H - and L -player, respectively, after HL or LH ; discounted $(w_{HH}^H, w_{HH}^L) \in \mathcal{E}^s(\delta)$ after HH ; and discounted $(w_{LL}^H, w_{LL}^L) \in \mathcal{E}^s(\delta)$ after LL .

These (common) expectations induce particular incentives for the two (symmetric) players governing their mixed action at time 0. Together with the continuation profile, this determines their expected repeated game payoffs at stage 0. With this in mind, we define a function $f : \mathcal{G} \rightarrow \mathcal{G}$, with \mathcal{G} denoting the set of all subsets of \mathcal{F} , which assigns to a candidate symmetric equilibrium set $\mathcal{E}^s \in \mathcal{G}$, used for continuation payoffs after HH and LL , the set of all possible equilibrium payoff-pairs at time 0, obtained by working through the appropriate incentives, as outlined above.

The set of symmetric equilibrium payoff-pairs $\mathcal{E}^s(\delta)$ must be a subset of the largest fixed point of f . The proof of the following proposition uses this fact in order to derive an upper bound $\bar{\mathcal{E}}^s(\delta)$ such that $\mathcal{E}^s(\delta) \subset \bar{\mathcal{E}}^s(\delta) \subset \mathcal{F}$.

Proposition 3. *In any symmetric equilibrium of the repeated allocation game characterized by $x = (1, 0)$, as the discount factor tends to one, no player can expect, conditional on being either the H - or L -player, a payoff that exceeds $\frac{3}{4}$. That is*

$$(w^H, w^L) \in \lim_{\delta \rightarrow 1} \mathcal{E}^s(\delta) \Rightarrow w^H, w^L \leq \frac{3}{4}$$

The proof is given in Appendix B. We provide here a brief sketch. We assume that the set of symmetric equilibrium payoffs satisfies the restriction that any individual payoff (to the H -player, or the L -player) cannot exceed a certain threshold \bar{w} . This assumption is certainly true for $\bar{w} = 1$. Now suppose we fix $\bar{w} \in (\frac{3}{4}, 1]$ and accordingly take continuations after HH and LL to satisfy that each payoff is less than or equal to \bar{w} , while, of course, the continuation after HL and LH is any pair in \mathcal{F} . We show that for every such profile of continuations the expected payoff pair to the H - and L -player at time 0 is strictly less than \bar{w} . This implies, in particular, that in order for a candidate set of payoff-pairs to be a fixed point of the mapping f , it must satisfy that no individual payoff exceeds $\frac{3}{4}$.

Note that Proposition 3 does not state that eventually, along the path of play, the payoffs to the H - and L -player cannot be highly asymmetric. Indeed, it is possible, from the point of view of the stage in which symmetries are broken, that the H -player receives a payoff of, e.g., $\frac{4}{5}$ and the L -player a payoff of $\frac{1}{5}$ (or anything nearby). What Proposition 3 does imply, though, is that, in this case, the expected time at which symmetry is broken, and those continuation payoffs are realized, is so far in the future (relative to δ) that the players' discounted payoffs viewed from time 0 are nowhere near efficient. This is true even if the discount factor is arbitrarily close to 1. The more patient players become, the longer it takes for symmetry to break, which, it turns out, is not compensated by the increased patience with which players view their payoffs.

Having shown that certain payoff-pairs are not possible to sustain with a symmetric equilibrium, we turn now to demonstrating that there are actually many payoff-pairs that *can* be supported, including asymmetric ones. In fact the next proposition completely characterizes the set of payoff-pairs sustainable in any stationary Markov equilibrium.

Proposition 4. *As δ tends to one, the set of payoff-pairs sustainable by symmetric Markov equilibria has Lebesgue-measure $\frac{1}{6}$, which is $\frac{1}{3}$ the measure of the feasible set.*

The proof is given in Appendix C. Proposition 3 and 4 are summarized in Figure 1. One can show, by example, that the set of payoff-pairs sustainable by all symmetric equilibria lies strictly between the bounds we provide. That is, one can sustain payoff-pairs (slightly) outside the lense depicted in Figure 1 via the use of non-stationary equilibria.




Figure 1: The lense represents the set of payoff-pairs sustained by symmetric Markov equilibria (Proposition 4). The corners represent payoff-pairs not achievable in any symmetric equilibrium (Proposition 3).

5 Implementing symmetric continuations

Corollary 1 says that, in repeated allocation games, the Pareto criterion identifies a particular unique payoff profile. Our view is that, at least in this context, strict Pareto optimality is a powerful basis on which to view such payoffs as focal and, therefore, to expect coordination on those payoffs. Nonetheless, if one is interested in a theory of play, which is to say a description of how to play the game, then there remains a multiplicity of strategies, in fact of symmetric equilibrium strategies, that deliver those payoffs. In this section we address the question as to how players might implement the ex-post symmetric payoffs for repeated 2-player allocation games.

Consider the stage in the repeated game at which symmetry is broken and players commence with an ex-post payoff-symmetric (but, necessarily, strategy-asymmetric) continuation strategy. Take any discount factor $\delta < 1$. It follows from Lemma 3.7.1 in Mailath and Samuelson (2007) originally due to Sorin (1986) that an exactly payoff-symmetric continuation can be constructed from an appropriate sequence of pure stage game action profiles, provided δ is large enough. Indeed it is easy to see that there are many such possible constructions (see, e.g., the proof of Proposition 3 in Bhaskar (2000)).

Normalize the time at which symmetry is broken to zero. Taking, for simplicity, and without loss of generality, the stage game of Section 4.1, one can describe any Pareto efficient continuation play by a sequence $y = (y_t)_{t=0}^{\infty} \in \{-1, 1\}^{\infty}$, where, without loss of generality, $y_0 = 1$.²⁷ The interpretation is that $y_t = 1$ corresponds to players using the same actions as in stage 0 (i.e., in which the H -player receives the high payoff), while $y_t = -1$ corresponds

²⁷We abuse notation by using y and z to denote continuation plays, rather than payoff profiles, in this section.

to players using the opposite actions (in which the L -player receives the high payoff). For any continuation y , define the normalized difference in payoffs between the two players by

$$\Delta^y(\delta) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t y_t.$$

We will also make use of the difference in payoffs corresponding to partial sequences: $\Delta^y(\delta|_T) = (1 - \delta) \sum_{t=0}^{T-1} \delta^t y_t$.

It is convenient to reproduce a version of Lemma 3.7.1 in Mailath and Samuelson (2007) here.

Lemma 1. *Let $\frac{1}{2} \leq \delta < 1$. There is a continuation play y such that $\Delta^y(\delta) = 0$.*

That is, for any $\delta \geq \frac{1}{2}$, one can find an exactly ex-post payoff-symmetric continuation play. However, the set of symmetric continuations is generally very sensitive to the discount factor δ . We are here interested in finding a particular sequence that is nearly symmetric for *all* large values of δ . With patient players, such a sequence represents the obvious way to play if one desires that the meta-norm of Pareto optimality be robust to small perturbations or uncertainty regarding the discount factor.

More precisely, this section accomplishes two goals. First, we prove that there is a unique such sequence. As it happens, this continuation is the well-known Thue-Morse sequence, which we denote throughout this section by z .²⁸ Specifically, we show that z satisfies the following property. For any k , and any sequence y with periodicity k , there exists a $\bar{\delta}$ such that for all $\bar{\delta} < \delta < 1$, $|\Delta^z(\delta)| < |\Delta^y(\delta)|$. Second, we provide a new characterization of z . For any δ , construct the continuation \hat{z} that, at every t , awards the payoff to the player who currently has the smaller total accumulated payoff.²⁹ That is, $\hat{z}_t = 1$ if $\Delta^{\hat{z}}(\delta|_t) \leq 0$ and otherwise $\hat{z}_t = -1$. We prove that the limit of sequences \hat{z} as δ tends to one is the Thue-Morse sequence, z . We begin with this latter result.

The Thue-Morse sequence, z , is defined as follows. Set $z_0 = 1$, and define the sequence recursively by $z_{2s} = z_s$ and $z_{2s+1} = -z_s$ for all s .³⁰

A simple lemma will prove useful.

Lemma 2. *Let z be the Thue-Morse sequence. Let $\mathbf{H}(t)$ be the number of 1's in the sequence z up to and including stage t . Define $\mathbf{L}(t)$, analogously, as the number of -1's in z through stage t . If t is odd then*

1. $z_t = -z_{t-1}$, and
2. $\mathbf{H}(t) = \mathbf{L}(t)$.

Proof: The first statement follows directly from the definition, $-z_{2s+1} = z_{2s} = z_s$. The second statement follows from the first (by induction). QED

²⁸See Thue (1906, 1912) and Morse (1921). Allouche and Shallit (1999) provide a useful discussion.

²⁹Bhaskar (2000) utilizes the same construction for a given, fixed, discount factor, showing that it is efficient in the infinitely repeated game. He does not, however, consider the limiting properties of this sequence as the discount factor tends to one. All of our results in this section are about this limiting case, where we identify the limit as the Thue-Morse sequence, and show that it has a robust optimality property for all high discount rates.

³⁰It is well known that this sequence can alternatively be constructed as follows. Set $z_0 = 1$. Proceed iteratively, at each step replacing every instance of 1 with $(1, -1)$, and replacing every instance of -1 with $(-1, 1)$. Another alternative construction is the following. Let us generally call $\{z_t\}_{t=0}^{2^k-1}$ the block of size 2^k . Let $\{-z_t\}_{t=0}^{2^k-1}$ denote the inverse of this block (i.e. 1's are replaced by -1's and vice versa). We have the block of size 1. The block of size 2 is then given by the block of size 1 followed by the inverse of the block of size 1. In general, the block of size 2^k is given by the block of size 2^{k-1} (which accounts for its first half) followed by the inverse of this very same block (which accounts for the second half). Yet another well-known equivalent characterization of the Thue-Morse sequence is the following. Let $z_t = 1$ whenever the binary expansion of t has an even number of 1's and let $z_t = -1$ otherwise.

We now show that the Thue-Morse sequence z has the property that, for sufficiently large δ , it awards the high payoff to the player who has the lower present discounted payoff.

Proposition 5. *For every t there exists a $\bar{\delta} < 1$ such that for all $\delta > \bar{\delta}$ the following is true. If $z_t = 1$ then $\Delta^z(\delta|_t) < 0$, and if $z_t = -1$ then $\Delta^z(\delta|_t) > 0$.*

Proof: The statement is obviously true for $t = 1$ and $t = 2$. Suppose now that $t \geq 3$ is odd. Then $\mathbf{L}(t-2) = \mathbf{H}(t-2)$ by the second part of Lemma 2. By the first part we then have $z_t = -z_{t-1}$. Suppose $z_{t-1} = 1$. Then $\mathbf{H}(t-1) = \mathbf{L}(t-1) + 1$ and, thus, there is a $\bar{\delta}$ such that for all $\delta > \bar{\delta}$ we have $\Delta^z(\delta|_t) > 0$. Given $z_t = 0$ the statement is true in this case. Now suppose that $z_{t-1} = -1$. Then $\mathbf{H}(t-1) = \mathbf{L}(t-1) - 1$ and, thus, there is a $\bar{\delta}$ such that for all $\delta > \bar{\delta}$ we have $\Delta^z(\delta|_t) < 0$. Given $z_t = 1$ the statement is true also in this case. This proves the statement for t odd.

We now turn to t even. Again, we know the statement is true for $t = 1$ and $t = 2$. Suppose the statement is true for all $\tau \leq t-1$. As t is even set $s = t/2$. We have

$$\begin{aligned} \sum_{i=0}^{2s-1} \delta^i z_i &= \sum_{i=0}^{s-1} \delta^{2i} z_{2i} + \sum_{i=0}^{s-1} \delta^{2i+1} z_{2i+1} \\ &= \sum_{i=0}^{s-1} (\delta^2)^i z_i - \delta \sum_{i=0}^{s-1} (\delta^2)^i z_i \\ &= (1 - \delta) \sum_{i=0}^{s-1} (\delta^2)^i z_i \end{aligned}$$

Since $s \leq t-1$, we know that if $z_s = 1$, then there is a $\bar{\delta}$ such that $\Delta^z(\delta^2|_s) < 0$ for all $\delta^2 > \bar{\delta}$ and so under the same condition, by the preceding development, $\Delta^z(\delta|_{2s}) < 0$. By definition, $z_t = z_{2s} = z_s = 1$, and the case is satisfied for all $\delta > \sqrt{\bar{\delta}}$. The argument is parallel for $z_s = -1$. QED

This result provides an economic interpretation of the Thue-Morse sequence. In particular, z is described by the limit of the sequence that, in every period, assigns the payoff to the currently disadvantaged player, as δ tends to one.

Before proceeding, recall that for any δ there exist sequences that are exactly payoff-symmetric. However, those sequences may become highly asymmetric for nearby values of δ . We now turn to showing that z has the desirable property of being very nearly payoff-symmetric for all δ close to one.

A sequence y has periodicity k if $y_{t+k} = y_t$ for all $t \geq 0$. We now prove the main result of this section.

Proposition 6. *For every k there exists a $\bar{\delta} < 1$ such that, for every sequence y with periodicity k , $|\Delta^z(\delta)| < |\Delta^y(\delta)|$ whenever $\delta > \bar{\delta}$.*

Proof: A sequence of periodicity k is characterized by its first k entries. Given its structure it is straightforward to express $\Delta^y(\delta)$ as a function of these first k entries. In fact,

$$\Delta^y(\delta) = \frac{1 - \delta}{1 - \delta^k} \sum_{t=0}^{k-1} \delta^t y_t.$$

Note that $\sum_{t=0}^{k-1} \delta^t y_t = \Delta^y(\delta|_k)$ is some polynomial in δ of degree $k-1$. We are interested in whether and, if so, at what rate, it converges to 0 as δ tends to 1. Let $A^0(\delta) = \Delta^y(\delta|_k)$. Suppose first that $A^0(1) \neq 0$. Then this sum, obviously, does not tend to 0. So suppose that $A^0(1) = 0$. Then $A^0(\delta)$ can be factored by $(1 - \delta)$. Thus, let $A^0(\delta) = (1 - \delta)A^1(\delta)$, where $A^1(\delta)$ is another polynomial in δ , but of degree $k-2$. We can, again, either have $A^1(1) = 0$ or not. In the latter case, the sum of interest then tends to zero at the same rate as $1 - \delta$.

In the former case, we can factor out another $1 - \delta$. Repeating this argument $k - 1$ times we obtain that there is a $\bar{\delta} < 1$ such that for all $\delta \geq \bar{\delta}$

$$|\Delta^y(\delta)| \geq \frac{(1 - \delta)^k}{1 - \delta^k}.$$

It remains to be shown that $\Delta^z(\delta)$ tends to zero faster than any such $\Delta^y(\delta)$. We have that $|\Delta^z(\delta)| = (1 - \delta) \prod_{j=0}^{\infty} (1 - \delta^{2^j})$.³¹ Set $\Delta_s^z(\delta) = \prod_{j=0}^{s-1} (1 - \delta^{2^j})$. Note that for any $j \geq 1$ the expression $(1 - \delta^{2^j})$ can be written as the product of $(1 - \delta^{2^{j-1}})$ and $(1 + \delta^{2^{j-1}})$. The former term can then, by the same argument, be factorized into another two such terms. Repeating, and applying this factorization to all terms in the above product we obtain an alternative representation of $\Delta_s^z(\delta)$, given by

$$\Delta_s^z(\delta) = (1 - \delta)^s \prod_{j=0}^{s-1} (1 + \delta^{2^j})^{s-j-1}.$$

Since, clearly, $\Delta_s^z(\delta) > \Delta_{s+1}^z(\delta)$ for all s , and $\Delta^z(\delta) = (1 - \delta) \lim_{s \rightarrow \infty} \Delta_s^z(\delta)$, it is the case that $\Delta^z(\delta) < (1 - \delta) \Delta_k^z(\delta)$.

Finally, we thus have

$$\begin{aligned} \Delta^z(\delta) &< (1 - \delta) \Delta_k^z(\delta) \\ &= (1 - \delta)^{k+1} \prod_{j=0}^{k-1} (1 + \delta^{2^j})^{k-j-1} \\ &< (1 - \delta)^{k+1} \prod_{j=0}^{k-1} 2^{k-j-1} \\ &= (1 - \delta)^{k+1} 2^{\sum_{j=0}^{k-1} k-j-1} \\ &= (1 - \delta)^{k+1} 2^{\frac{k(k-1)}{2}}. \end{aligned}$$

Thus, as δ tends to 1, $\Delta^z(\delta)$ tends to zero at least an order faster than $\Delta^y(\delta)$ when y has a given, but arbitrary, periodicity k . This completes the proof. QED

One can extend the characterization of z to n -player allocation games, with $n \geq 3$. The most obvious way to do so is to consider, at any stage t in the continuation play, the ranking of players, π^t , defined by increasing order of their present discounted payoffs in the repeated game. The play at stage t , then, is defined by $a_{\pi^t(i)}^t = i$, i.e., the action of the player with the i -th lowest payoff at stage t is the one that pays the i -th largest amount according to x .

Consider high values of δ . For any allocation game defined by a given vector x , this sequence is well-defined.³² For such δ , this continuation play is ex-post payoff-symmetric. It also has the property, as in the 2-player case, that in every successive block of n periods, each player receives each of the payoffs x_a exactly once. The exact sequence to which this play converges, as δ tends to one, will generally depend on x . We conjecture that for any repeated allocation game, that limit sequence retains the property ascribed to the Thue-Morse sequence by Proposition 6, i.e., that it is more payoff-symmetric than any periodic sequence for sufficiently high discount factors.

³¹Calling $\{z_t\}_{t=0}^{2^j-1}$ the j -th block, this can be verified by observing that the difference in payoffs after the first block is $\Delta_0^z(\delta) = 1$. After the second block it is $\Delta_1^z(\delta) = (1 - \delta) \Delta_0^z(\delta)$, after the third $\Delta_2^z(\delta) = (1 - \delta^2) \Delta_1^z(\delta)$, and generally after the $j + 1$ -st it is $\Delta_j^z(\delta) = (1 - \delta^{2^{j-1}}) \Delta_{j-1}^z(\delta)$. Finally the desired difference is $\Delta^z(\delta) = (1 - \delta) \lim_{j \rightarrow \infty} \Delta_j^z(\delta)$, which can be expressed as the infinite product given above.

³²It is well defined generically, the only concern being cases in which two players have equal present discounted payoffs.

6 Simplicity

We now consider the matter of the complexity of the strategy that is to serve as the focal point of the repeated allocation game.

According to Kalai's (1990) survey on bounded rationality and complexity in repeated games, Aumann (1981) was the first to point to viewing repeated game strategies as automata as an aid to measuring complexity of those strategies. The most common measure of complexity of a given repeated game strategy was then formally provided (simultaneously) by Ben Porath (1986), Neyman (1985), and Rubinstein (1986) as *state-complexity*, which is given as the minimal number of states of any (finite) automaton that encodes the given repeated game strategy. The higher this number the more complex is a strategy. Kalai and Stanford (1988) then showed that this notion of complexity of a repeated game strategy is equivalent to the number of continuation strategies this repeated game strategy can generate. Thus, state-complexity can be identified without reference to automata.

We shall use (low) complexity very differently from the work reviewed by Kalai (1990). In Ben Porath (1986) and Neyman (1985), among others, players are allowed to use only strategies of bounded complexity. In Rubinstein (1986) and the literature following it, as reviewed in Chatterjee and Sabourian (2009), players have, in addition to preferences over outcomes, also a preference in favor of less complex strategies. We here appeal to low complexity simply as a selection device (choosing from all complexity-unrestricted symmetric equilibria) to complete our description of a focal point in repeated allocation games and, thus, follow a third strand of the complexity literature which was initiated by Baron and Kalai (1993).

We thus define the complexity of a strategy σ as the smallest number of states in any automaton that implements this strategy.

Proposition 7. *Consider any repeated n -player allocation game Γ . Let σ denote a strategy with the property that the symmetric strategy profile $\{\sigma\}^n$ is efficient in the limit as δ tends to one. Let $(\mathcal{W}, w^0, f, \tau)$ be an automaton representation of σ . Then $|\mathcal{W}| \geq n + 1$.*

Proof: Note first that there must be a state $w \in W$ such that $f(w)$ is totally mixed over A . Note also that in order for this automaton to eventually lead to an efficient continuation, it must be that all actions must be played purely after some history. I.e. for every action $a \in A$ we must have a state w_a such that $f(w_a)$ attaches probability 1 to a single action. QED

There are, of course, less complex automata. For instance consider the automaton with just a single state describing to play action 1. The symmetric profile of such a strategy is, in fact, an equilibrium (unless $n = 2$ and $x \gg 0$) and it yields a payoff of zero. There are also less complex automata that implement strategies with strictly positive payoff profiles. For example, the single-state automaton that prescribes the totally mixed stage game symmetric Nash mixture at every history constitutes a symmetric equilibrium. What Proposition 7 says is that all symmetric strategy profiles (whether or not they are equilibria) with fewer than $n + 1$ states are necessarily inefficient for patient players.

As the discount factor tends to one, how do players select a focal strategy that implements the efficient payoff profile? The answer that we provide in Section 5 results in the Thue-Morse sequence. The Thue-Morse sequence can be represented by a very simple Turing Machine. We need one memory cell with an initial element, one function of the memory to $\{0, 1\}$, and one function of the memory to itself that tells us how to update the memory.³³

³³For instance, we can use an initial memory $m_0 = 0$; output function $f : \mathbb{N}_0 \rightarrow \{0, 1\}$ given by $f(m) = 1$ if m 's binary expansion has an even number of zero's (and $f(m) = 0$ otherwise); and memory update function $g : \mathbb{N}_0 \rightarrow \mathbb{N}_0$ given by $g(m) = m + 1$. Alternatively, and close to the motivation of Section 5, we can use initial memory $m_0 = 1$; output function $f : \mathbb{R} \rightarrow \{0, 1\}$ given by $f(m) = 1$ if $m < 0$ (and $f(m) = 0$ otherwise); and memory update function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $g(m) = g(m) + D$, where D is the current payoff difference.

But the Thue-Morse sequence cannot be represented by a finite automaton. Supposing that complexity, in players' minds, is well-captured by state-complexity of minimal automaton representation, we now ask the following question. Among symmetric equilibria that deliver efficient payoffs as δ tends to one, which strategy is simplest and, thus, a candidate focal point of the game?

Fix an allocation game Γ . Fix an order over, or labeling of, actions $g : A \rightarrow \{1, \dots, n\}$ that is surjective.

Definition 4. A rotation scheme for Γ with respect to order g is a strategy with an automaton representation $(\mathcal{W}, w^0, f, \tau)$ with state space $\mathcal{W} = R \cup A$, initial state $w^0 = R$, output function f with $f(R)$ specifying a totally mixed distribution over A , and $f(a) = a$ for each $a \in A$, and transition function τ that produces $\tau(R, a, n) = R$ (for the case when symmetry is not broken), and $\tau(R, a, c) = \tau(a, \cdot, \cdot) = g^{-1}(g(a) + 1 \bmod n)$ (otherwise), where a is the realized action of the player.

In words, a rotation scheme randomizes with a fixed distribution until symmetry is broken, and then rotates among the pure actions according to g in the continuation. Note that a rotation scheme achieves the lower bound on state complexity identified in Proposition 7. Generically, the simplest efficient symmetric equilibria for patient players are rotation schemes of Γ . We focus on allocation games with the following property.

Definition 5. A vector of payoffs x for an n -player allocation game is strongly distinct if, given two non-identical sets of non-negative integers $\{\alpha_a\}_{a \in A}$ and $\{\beta_a\}_{a \in A}$, we have $\sum_{a \in A} \alpha_a x_a \neq \sum_{a \in A} \beta_a x_a$.

The following result formalizes the characterization.

Proposition 8. Consider an n -player allocation game with strongly distinct payoff-vector x . Consider the set of strategies Σ^* such that their associated symmetric strategy profiles constitute efficient equilibria in the limit as δ tends to one. If $\sigma \in \Sigma^*$ has an automaton representation with state-complexity $n + 1$, then σ is a rotation scheme.

Proof: Efficiency of the strategy implies that there must be at least one state in which the play is totally mixed. By Corollary 1 an efficient equilibrium must be ex-post payoff-symmetric. Since x is strongly distinct, as δ tends to one, the requirement that every player achieve the same discounted payoff implies that every player must receive each of the $\{x_a\}_{a \in A}$ payoffs infinitely often. The unique simplest way to do this is by rotating through actions according to a surjective rotation g . QED

Proposition 8 identifies as rotation schemes those strategies that are the simplest efficient symmetric equilibria for patient players. Uniqueness, though, is determined only up to the order, g . For a rotation scheme to qualify as a focal point of the repeated allocation game, it must be that in players' minds there is a conventional order over the actions.

That order must be defined with reference only to the game itself, which is to say that g should depend only on x , the vector of payoffs of the allocation game. While there are many ways to construct such a function, the notion of simplicity suggests a particularly natural construction, which might be called "increasing." That is, upon successful breaking of symmetry, players coordinate by rotating through the payoffs of x in increasing order (mod n).³⁴ Once "increasing" is determined to be the conventional ordering, Proposition 8 identifies a unique way to play the game. While it is not exactly efficient for $\delta < 1$, it has the desirable property of being most simple among all nearly efficient continuation protocols.

³⁴While "increasing" is, in our view, the most natural candidate for a simple rotation scheme, it may be argued that "decreasing" is equally as natural a choice. As we demonstrate below, in our laboratory experiments, whenever a rotation scheme is adopted, which is very often, it is always "increasing".

7 Experimental Evidence

The theory of focal points that we have developed makes a number of testable predictions for play in repeated allocation games. In order to assess the empirical content of the theory, we report data from a battery of laboratory experiments implementing these games.

7.1 Description of Experiments

The experimental sessions were conducted at the Social Science Experimental Laboratory (SSEL) at Caltech between February and November, 2010. We recruited undergraduate subjects who had no previous experience with related experiments. For each session, subjects entered the laboratory and were randomly assigned to private computer terminals separated by dividers. All interactions among subjects were computerized, using an extension of the open source software package Multistage Games.³⁵ Instructions were read aloud for everyone to hear. The exchange rate from experimental points to dollars varied across sessions so that the predicted earnings of each subject was \$15 – 25, inclusive of a \$10 show-up payment. Each experimental session lasted approximately one hour, including instructions (which are available as supplementary material) and payment.

Each experiment consisted of ten matches. In each match subjects were (uniformly and independently) randomly allocated into groups to play the repeated allocation game. The stage games were implemented as follows. In the 2-player games there were two actions, chosen with on-screen buttons labeled L and H . In the 3-player games there were three buttons, labeled E , W , and H . In either case, the experimental screen displayed these buttons in order from left to right. Unless the action profile was a permutation of the set of pure actions, all payoffs were zero. Otherwise, positive payoffs were awarded.

We varied several aspects of the repeated allocation game across sessions. The main design variables are the number of players, the stage game payoffs, the stopping rule, discounting, and feedback. These parameters are summarized in Table 1. We turn now to a detailed description of these treatments.

The vector of coordinated payoffs, x , takes values $(30, 10)$, $(30, 1)$, $(30, 20, 10)$, $(30, 20, 1)$, and $(30, 5, 1)$, denominated in experimental points, as shown in the third column of Table 1. The main purpose of varying x is to vary the magnitude of possible asymmetry in payoffs. In sessions 4 and 5, payoffs (marked by an asterisk in Table 1) were discounted with a factor $\delta = 11/12$ across rounds. In all other sessions payoffs were not explicitly discounted.

The stopping rule was either Fixed or Random in each session. In the Fixed Stopping condition, each match consisted of twenty rounds. In the Random Stopping condition, a pair of dice were rolled at the end of each round. If the sum of the dice was less than or equal to three (probability $3/36$), the match was terminated; otherwise it continued to a next round. Notice that the implied discount factor under Random Stopping is analogous to the rate at which payoffs are discounted in sessions 4 and 5. As we discuss below, for a given continuation play, the predicted mixing probabilities before symmetry is broken depends on the stopping rule.

Finally, we varied the feedback (monitoring) across sessions. The main feedback treatment is called Mixed. Under Mixed Feedback, from the onset of a match, subjects observe only their own payoff after each stage game, until symmetry is broken. In every stage of the continuation in which symmetry is broken, subjects observe both their own payoff and the actions taken by the other players. The second feedback treatment is called Payoff. Under Payoff Feedback, subjects observe only their own payoff at every stage of the game.

In the 2-player games, the treatments are equivalent, since one's own payoff implies a certain action for the opponent. But Payoff Feedback provides strictly less information than

³⁵Documentation and instructions for downloading the software can be found at <http://multistage.ssel.caltech.edu>.

Session	Players	Payoffs	Stopping Rule	Feedback	Subjects
1	2	(30, 10)	Random	Mixed	10
2	2	(30, 1)	Random	Mixed	10
3	2	(30, 10)	Fixed	Mixed	12
4	2	(30, 10)*	Fixed	Mixed	16
5	2	(30, 10)*	Fixed	Difference	14
6	3	(30, 20, 10)	Random	Payoff	12
7	3	(30, 20, 1)	Random	Payoff	15
8	3	(30, 5, 1)	Random	Payoff	15
9	3	(30, 20, 10)	Random	Mixed	15
10	3	(30, 5, 1)	Random	Mixed	15
11	3	(30, 20, 10)	Fixed	Mixed	15
12	3	(30, 20, 10)	Fixed	Mixed	15
13	3	(30, 5, 1)	Fixed	Mixed	15

Table 1: Summary of experiment sessions.

Mixed Feedback in the 3-player games. Notice though that the distinction is irrelevant theoretically, since the breaking of symmetry is, on its own, enough information to allow unrestricted joint play. Nevertheless, we shall see that the distinction is empirically relevant.

There is a third feedback treatment called Difference, used in a single 2-player session. Under Difference Feedback, subjects were implicitly given the same information as in Mixed Feedback, but it was framed differently. The summary statistic that subjects were given on-screen was the difference between their current cumulative payoff in the match and that of their opponent. The idea behind Difference Feedback was to encourage continuation play that resembled the Thue-Morse sequence.

Under all feedback treatments, the computer screen contained a history panel giving feedback from previous stages in the match, according to the specific treatment for that session. At the end of every session, points were summed from the subjects' earnings across the ten matches and converted to dollars. Subjects were paid privately in cash at the end of the session.

7.2 Symmetry in the laboratory

Before proceeding with the analysis of the data, it is important to comment on the relationship between the theory of focal points and the conditions in the laboratory.

The informational features of the laboratory setting we implement ensure that the restrictions of symmetric play are appropriate for a strategic analysis of observed behavior. By design, all subjects in a given group are in an exactly symmetric position with each other at the start of each match. The matching process is random and anonymous, so subjects have no information about who their opponents are. They, therefore, cannot use any norms that operate on such distinctions between players. Players continue to be in a symmetric position until a realization of positive payoffs in the group. At this point, under all of our feedback treatments, it is common knowledge among the group that there is a particular subject who got each available distinct positive payoff. For the remainder of the match, no two players are symmetric, and symmetry thus places no restriction on the continuation play.

Once symmetry is broken, there is of course the matter of what, if any, continuation play is focal. If there is any lack of common knowledge of the focal play, then further information could be useful in coordinating play. As we will see, the Mixed Feedback treatment has a

powerful effect in this direction.

Of course, it is possible that subjects in exactly symmetric positions will choose different strategies. This could arise for a variety of behavioral reasons, potentially leading to non-equilibrium play. Most important from our point of view, however, is that all rational calculations are indeed subject to the restrictions of symmetry. In particular, there is no public or private information available to subjects that would allow them to coordinate on asymmetric play when symmetries are present. Even if such a signal were available, it would be extremely difficult to use as a coordination device unless there was a norm associated with the signal that served as a way to commonly interpret it.

Even so, the theory that we develop above is not directly applicable in all respects to the setting of the experiments. The analysis above proves results for infinitely repeated allocation games with discounting, focusing on the limiting results as players become perfectly patient. These conditions are not feasible to implement in the laboratory. Most of the strategic effects we are interested in, however, do not change dramatically when moving from the theoretical setting to the laboratory.

In particular, the theoretical argument for focusing on Pareto efficient continuations once symmetries are broken is equally compelling in the finite horizon case or the case of impatient players as in the infinitely repeated version with a high discount factor. Similarly, the observation that the closer an equilibrium is to delivering ex-post symmetric payoffs, the more efficient will be the mixing probabilities before symmetries are broken, remains true in finite horizon games.

In general, the equilibrium set of a finitely repeated game may be very different from its infinitely repeated counterpart. But the effects in our setting are relatively minor. This is because, for the allocation games we study, in every symmetric equilibrium, once symmetries are broken, the continuation play in every stage of the infinitely repeated game is a stage game Nash equilibrium. There is no unraveling of candidate equilibria due to backwards induction. In addition, there are some new theoretical predictions that arise in the case of finite horizon games, and these are readily testable with our data.

7.3 Efficiency of continuation play

We assess first our contention that continuation play following the breaking of symmetry will be Pareto efficient. Our hypothesis is that in allocation games Pareto efficiency is likely to be a salient meta-norm. This entails that once positive payoffs are achieved in a match, the play in every subsequent round will be coordinated and also result in positive payoffs.

Pooling all of the 2-player sessions, we find that in all of the rounds that occur after symmetry has been broken, 93% (4082/4398) successfully achieve coordination. Given that there are many feasible efficient continuations, and achieving one of them requires coordination between the players, we view this as very strong evidence in favor of our hypothesis.

Turning to the 3-player sessions, we find that the feedback treatment is very important. The reason is that coordination is a fundamentally harder objective to achieve in 3-player games. Notice that only 6 of the 27 (22%) possible action profiles result in positive payoffs. Pooling the sessions from the baseline treatment, in which subjects receive Payoff Feedback, 34% (300/876) of rounds in which symmetry was broken achieve coordination.

While this is significantly better than random, part of the reason that play was not more efficient is due to miscoordination between players on how to proceed in the continuation play. The Mixed Feedback treatment is meant to address this issue, whereby players can see what actions their opponents are choosing when coordination fails, and use this information to help coordinate in future rounds. Pooling across the Mixed Feedback sessions, 68% (2417/3554) of rounds in which symmetry was broken achieve coordination.

7.4 Ex-post payoff symmetry

As discussed above, given that players start each match under symmetry, it must be that the ex-ante expected payoffs are the same for all players. But different strategies have different ex-post payoff asymmetries. Continuation plays that are closer to ex-post payoff-symmetric induce mixing probabilities that are closer to uniform. Such play breaks symmetries faster, resulting in Pareto superior payoffs ex-ante. Based on this reasoning, the meta-norm of Pareto efficiency suggests not only that Pareto efficient continuations will be used. Beyond that, nearly payoff-symmetric continuations should be used.

To assess the extent to which this is true in our data, we examine how close payoffs are to being ex-post symmetric. We denote by $\pi = (\pi_1, \dots, \pi_n)$ the vector of realized payoffs for a repeated allocation game, with $\pi_1 \geq \dots \geq \pi_n$. Define for a given repeated allocation game

$$S = \frac{(\pi_1 - \pi_n)/\pi_n}{(x_1 - x_n)/x_n},$$

which measures a normalized percentage difference in ex-post payoffs between the highest and lowest paid player. $S = 0$ when payoffs are ex-post symmetric, and S is normalized to have maximal value at unity, which corresponds to the maximally asymmetric continuation play in which every subject takes the same action in every round following the breaking of symmetry. We aggregate S across groups and matches within a session by taking a weighted average across observations, with weights given by the total payoff of the group in the repeated game.

The mean value of S across the five 2-player sessions is 0.05. In the 3-player games we find again that the feedback treatment has a large effect. Under Payoff Feedback, in which it is harder to play complex continuations due to lack of information about the action choices of others, subjects more frequently played the simplest continuation in which players use the same actions in subsequent rounds, delivering very asymmetric payoffs. This is captured by the large average figure of $S = 0.51$ for the three Payoff Feedback sessions. In the remaining 3-players sessions, which used Mixed Feedback, the average value of S is 0.13.

As a robustness check we consider an alternative measure of payoff asymmetry given by

$$S' = \frac{(\pi_1 - \pi_n)/R}{x_1 - x_n},$$

where R is the number of rounds in the continuation game from the point at which symmetry is broken. One difference between S' and S is that they treat differently the normalization between games with a very low value of x_n and those with more equal payoffs. The results that we obtain, though, are qualitatively similar using both measures. These results, along with those on the efficiency of continuation play, are depicted at the session level in Table 2.

All of these results show that subjects clearly manage to use sufficiently complex continuations in order to reduce ex-post asymmetry in payoffs. This is true even in the 3-player Payoff Feedback sessions, and much more so in the other sessions. Our theory explains this play as arising from a norm of ex-ante Pareto efficiency.

7.5 Continuation strategies: 2-player games

Thus far we have provided evidence that, as predicted, subjects in our experiments play repeated allocation games in such a way as to achieve relatively ex-post payoff-symmetric outcomes. We turn now to analyzing which strategies are used to deliver such payoffs. A useful starting point for analyzing the continuation play after symmetry is first broken is to determine which continuations are most payoff-symmetric in the experimental treatments. The answer to this question depends on the design variables of the experiment.

Session	Efficiency of Continuation Play				Ex-post Payoff Symmetry			
	# Rounds	# Coord.	Ratio	Totals	S	Avg.	S'	Avg.
1	223	215	96%		0.084		0.101	
2	446	437	98%		0.009		0.044	
3	1088	995	91%	93%	0.040	0.050	0.043	0.050
4	1412	1369	97%		0.029		0.024	
5	1229	1066	87%		0.087		0.040	
6	356	134	38%		0.662		0.539	
7	289	92	32%	34%	0.604	0.591	0.447	0.435
8	231	74	32%		0.508		0.318	
9	844	610	72%		0.276		0.218	
10	421	274	65%		0.050		0.088	
11	753	456	61%	68%	0.101	0.134	0.072	0.121
12	778	563	72%		0.223		0.173	
13	758	514	68%		0.021		0.052	

Table 2: Efficiency and ex-post payoff-symmetry of continuation play.

We consider first the 2-player games under Random Stopping (sessions 1 and 2). Payoff symmetry requires that the expected continuation payoffs to playing H and L , conditional on symmetry being broken, are equal. This is equivalent to the requirement that, in the infinitely repeated game with discount factor equal to the probability of continuation under Random Stopping, the (non-random) continuation payoffs to playing H and L , conditional on symmetry being broken, are equal.³⁶ Recall that there are many such continuations that deliver exactly symmetric payoffs in the infinitely repeated game, provided δ is large enough, which holds in our experiments. One example is the construction of awarding the high payoff to the player with the smaller current cumulative payoff. A different construction that is also exactly symmetric works by awarding the high payoff to the player who won the high payoff at the previous stage, provided that his new normalized discounted payoff would not exceed $(x_1 + x_2)/2$, in which case the payoff is awarded to the opponent.

However both of these continuations, and every other exactly egalitarian continuation at $\delta = 11/12$, are very complex in the sense of automaton state complexity.³⁷ The multiplicity of ex-post symmetric continuations, together with the complexity of each one of them, makes such continuations unlikely candidates for observed behavior. There are much simpler strategies that are very nearly payoff-symmetric. The simplest among these are rotation schemes or, in the case of two players, turn-taking.³⁸

We consider next the 2-player games under Fixed Stopping without discounting (session 3). All two player games have a High payoff (H) and a Low payoff (L). Any continuation play that splits the H and L payoffs equally between the two players, and only these continuations, are most payoff-symmetric. It does not matter in what order the H and L payoffs are awarded. When there are an even number of rounds left at the moment symmetries are broken, each player must receive the H payoff in exactly half of the remaining rounds. When there are an odd number of rounds remaining, one player must receive H exactly one more time than the

³⁶As described above, one could consider strong ex-post payoff-symmetry under Random Stopping, which is a strictly stronger requirement.

³⁷To illustrate, the first 20 elements of the sequences are, respectively, $(1, 0, 0, 1, 0, 1, 1, 0, 0, 1, 1, 0, 1, 0, 0, 1, 0, 1, 1, 0, \dots)$ and $(1, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, \dots)$.

³⁸For $\delta = 11/12$, in all of the games we study, the ex-ante efficiency loss in equilibrium associated with rotation, relative to the exactly payoff-symmetric strategies, is never more than 1%; see Table 5.

other player.

Thus, as under Random Stopping, many continuation profiles deliver payoffs that are maximally ex-post symmetric. Efficiently implementing a payoff-symmetric outcome, though, requires the players to coordinate on a particular continuation play from the many acceptable ones. Our analysis above suggests that, among all such continuation plays, simplicity is the natural focal criterion for selecting a continuation play in this setting. As proven above, rotation is the simplest such continuation, in the sense of minimal state complexity of an automaton that implements the strategy.³⁹

The remaining case for 2-players games is Fixed Stopping with discounting (sessions 4 and 5). Here, the maximally payoff-symmetric continuation play is unique, and it depends on the round in which symmetry is broken. These continuations are presented in Table 3.⁴⁰ It is obvious that this continuation protocol is highly complex and changes unpredictably with the round at which symmetry is broken. We thus predict that a much simpler but nearly payoff-symmetric continuation play will be used instead. Rotation is, again, one obvious candidate. Session 5 was designed to encourage the continuation play in which the player with the lower current cumulative payoff plays H, in order to assess how robust the convention of rotating is in these experiments.

Summarizing, our hypothesis is that rotation schemes will be observed frequently in sessions 1-4. In session 5 we expect to observe either rotation or the awarding of the high payoff to the currently disadvantaged player. We examine two summary statistics regarding rotation schemes. The first is very demanding. We look at every observation in which a group broke symmetry with at least one round of continuation play remaining. Most of these observations, of course, have many rounds remaining. Among all such groups, we count the number of groups that immediately enter a rotation scheme and maintain it perfectly until the end of the match.

The findings are summarized in Table 4. Pooling all of the 2-player sessions, we find that 81% of these groups play a perfect rotation scheme in the continuation play. While already high, this statistic omits many groups who have nearly perfect rotation, or who have perfect rotation from some point on that is strictly after the round when symmetry was broken. To account for such observations, we look also at the frequency of groups who end the match with $n = 2$ rounds of rotation. In this case the observed frequency of success raises to 89%. We thus take this as strong evidence that rotation schemes are the dominant choice of continuation play. Even in session 5 81% of groups end with rotation, and 74% of groups rotate perfectly upon the first breaking of symmetry, despite the framing of the feedback. Simplicity, paired with Pareto efficiency, is indeed a powerful criterion for focality.⁴¹

7.6 Continuation strategies: 3-player games

The main thrust of our argument and the conclusions from the 2-player games hold also in 3-player games. We discuss them separately only because there is an additional subtlety that arises. That subtlety is that the rotation scheme is not unique when $n > 2$. As discussed in Section 6, there is one rotation scheme for every ordering over actions. One particular such rotation scheme must be coordinated on in order to achieve positive payoffs.

³⁹The proof of Proposition 8 is easily adapted to the finite horizon case. We remark that, in the case of non-discounted payoffs, rotation also coincides with a version of awarding the high payoff to the player with the lower cumulative payoff, in which ties are always awarded to the same player.

⁴⁰We remark that this finding contradicts Proposition 4 in Bhaskar (2000), which claims incorrectly that for any finite horizon and any discount factor, the convention of awarding the higher payoff to the currently disadvantaged player is most egalitarian. See Kuzmics and Rogers (2012) for details.

⁴¹These criteria are not special to the laboratory, as suggested by the example of rotation schemes in pickup sports games mentioned in footnote ??.

Round	Rounds remaining in continuation play upon the breaking of symmetry																			
	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	0	0	0	1	0	0	0	1	0	1	0	0	0	1	0	1	0	0	0	0
3	0	0	0	0	1	0	0	1	0	0	1	1	0	1	0	0	0	0	0	0
4	0	1	0	1	1	0	1	0	1	1	0	1	1	0	1	0	1			
5	0	0	1	1	0	1	1	0	0	0	0	0	0	0	0	0				
6	1	1	0	0	0	0	0	0	1	0	1	1	1	0	1					
7	1	1	1	0	0	1	0	0	1	1	0	0	1	0						
8	1	0	1	0	1	1	1	0	0	0	1	0	0							
9	0	1	0	1	0	1	1	1	0	0	0	0								
10	0	0	1	0	0	0	0	1	1	0	1									
11	1	1	0	1	0	0	0	0	1	1										
12	1	0	0	0	1	1	1	0	0											
13	1	1	1	0	1	1	0	1												
14	1	1	1	1	1	0	1													
15	0	0	1	0	1	1														
16	0	0	0	0	0															
17	1	1	1	0																
18	1	0	1																	
19	1	0																		
20	0																			

Table 3: The columns represent the maximally payoff-symmetric continuation for 2-player games under Fixed Stopping with payoffs discounted at rate $\delta = 11/12$. “1” represents the H-player playing H and “0” represents the H-player playing L.

Session	# Obs.	Perfect Rotation or Flipping			Ending Rotation or Flipping		
		# Rotate/Flip	Ratio	Totals	# Rotate/Flip	Ratio	Totals
1	44	39	89%		40	91%	
2	45	41	91%		43	96%	
3	60	41	68%	81%	51	85%	89%
4	80	70	88%		74	93%	
5	69	51	74%		56	81%	
6	22	0	0%		2	9%	
7	28	1	4%	3%	1	4%	5%
8	25	1	4%		1	4%	
9	43	10	23%		19	44%	
10	38	10	26%		16	42%	
11	49	12	24%	30%	26	53%	52%
12	50	13	26%		29	58%	
13	50	23	46%		29	58%	

Table 4: Frequency of rotation schemes.




Figure 2: Examples of play from three matches, showing a rotation scheme (left), flipping (center) and the simplest, but payoff-asymmetric, play of remaining with the same actions (right). The data is from Session 11, match 3 and group 4 (left), match 4 and group 3 (center), and match 4 and group 5 (right).

In every session and every group within each session that plays a rotation scheme, the observed order over actions is that of increasing payoffs. The unanimity of this convention is striking.⁴²

Four of the sessions (6, 9, 11, and 12) have payoffs given by $x = (30, 20, 10)$. This is a knife-edge case that violates the generic condition of strongly distinct payoffs. As a result, in these sessions it is not true that rotation schemes are the uniquely simplest strategies that are ex-post payoff-symmetric as δ tends to one. Indeed, there is also the “flipping” convention in which $30 \rightarrow 10 \rightarrow 30$ and $20 \rightarrow 20$. In fact, there is a sense in which this scheme may be considered simpler than rotation. There is an automaton representation of flipping that can be reduced to two states after a finite amount of time, which is not true of a rotation scheme.

To illustrate the typical patterns of play in our data we refer to Figure 2. Each panel depicts the play of one group in one match from one session (Session 11). The horizontal axis codes the rounds within the given match, while the vertical axis codes the payoff of subjects in each round. Subjects are color coded for clarity. The left panel shows a rotation scheme, the center panel shows an example of flipping, and the right panel shows a group that, upon the breaking of symmetry, played the same actions in all remaining stages of the game, resulting in asymmetric payoffs. Notice that in this final case, it happened to take longer for subjects to break symmetry.

As noted above, the 3-player sessions with Payoff Feedback (sessions 6-8) suffer from an inability of subjects to effectively coordinate their continuation play. As a result, the proportion of groups that end the match with a round of rotation or flipping is very small – 5% in aggregate. The groups in the remaining sessions (9-12) average 52%, with 30% rotating or flipping perfectly starting when symmetry is broken. Again, as perfect rotation is demanding requirement, and all the more so in 3-player games in which there are multiple ways to rotate or flip, we view this as a very positive result.

The groups that coordinate efficiently in sessions with $x = (30, 20, 10)$ use exclusively the flipping scheme and never a rotation scheme. Moreover, these sessions have weakly higher rates of perfect flipping than the rotation rates in sessions with strongly distinct payoffs. Both of these findings are consistent with our observation that flipping may be considered more simple than rotation.

⁴²We should point out that the order of the action buttons on subjects’ screens from left to right coincided with increasing payoffs. So it is possible that the convention was “left to right.” We cannot identify which of these conventions drives our finding.

	Payoff Vector				
	(30, 10)	(30, 1)	(30, 20, 10)	(30, 20, 1)	(30, 5, 1)
Strategy	(.569, .413)	(.625, .375)	(.349, .333, .317)	(.347, .356, .297)	(.378, .339, .283)
Value	18.41	14.16	11.99	10.16	7.13
Efficiency	99.7%	99.0%	99.9%	99.6%	99.0%

Table 5: Mixing probabilities before symmetry is broken under Random Stopping. Mixing probabilities are then used to compute the value of the game, which is also listed as a proportion of the value obtained under optimal uniform mixing.

7.7 Breaking symmetry and equilibrium mixing: Theoretical Predictions

Having presented the findings on continuation play from the stage in which symmetry is broken, we turn now to analyzing play in stages where symmetry is present. Our analysis provides an explicit link between the continuation play and mixing probabilities in the presence of symmetry. Indeed, one of the main points of the analysis is that continuation strategies that deliver more egalitarian payoffs have mixing probabilities that are closer to uniform, and hence are more efficient. Before presenting the empirical findings, we first extend the theory to the games that are implemented in the lab.

Given the predominance of rotation and, in the case of $x = (30, 20, 10)$, flipping, observed in the data, we conduct our main analysis of mixing probabilities under the assumption that all players anticipate this continuation play upon the initial breaking of symmetry.

We first discuss Random Stopping. Mixing probabilities in the symmetric equilibrium depend on the continuation payoffs and the discount factor. We describe the computation of equilibrium mixing for the case of the 2-player games; the three player games work similarly. The utility of action H is computed as

$$u(H) = \alpha v + (1 - \alpha) \frac{x_1 + \delta x_2}{1 - \delta^2},$$

where α is the probability the opponent plays H and v is the value of the game. The utility of L, $u(L)$, is computed similarly. The symmetric equilibrium values of α and v are those which solve the system $u(H) = u(L) = v$. Under a rotation or flipping scheme, equilibrium mixing probabilities are stationary. Finally, the (normalized discounted) value of the game can be compared to the value that would obtain under optimal uniform mixing. Table 5 summarizes these computations applied to our sessions.

In the Fixed Stopping treatment, equilibrium mixing is not stationary. The mixing probabilities are easily computed, though, via backward induction. We again describe the procedure for the two player games; the three player games work similarly. In the final round, round twenty, the mixing probabilities are simply those of the (Pareto efficient) symmetric Nash equilibrium. This equilibrium has associated with it a value, $v(20)$. In the penultimate round, the utility of action H is computed as

$$u(H|19) = \alpha v(20) + (1 - \alpha)(x_1 + x_2),$$

where α is the probability the opponent chooses H. This is because when playing H, with probability α symmetry is not broken and the total value of the game reflects the value at the last stage, $v(20)$, while with probability $(1 - \alpha)$ symmetry is broken and a payoff of x_1 is obtained in stage 19 and, given the rotation scheme, a payoff of x_2 is obtained in the last stage. The payoff to action L is computed similarly, and the equilibrium mixing probability is that which equates the expected utilities to the two actions. The value of the game at

Round	Payoff vector		
	(30, 10)	(30, 20, 10)	(30, 5, 1)
1	uniform	uniform	(.357, .370, .273)
2	(.622, .378)	(.354, .333, .313)	(.394, .310, .296)
3	uniform	uniform	uniform
4	(.622, .378)	(.355, .333, .312)	(.358, .372, .270)
5	uniform	uniform	(.397, .308, .295)
6	(.623, .377)	(.356, .333, .311)	uniform
7	uniform	uniform	(.359, .374, .267)
8	(.623, .377)	(.357, .333, .310)	(.402, .307, .291)
9	uniform	uniform	uniform
10	(.623, .377)	(.359, .333, .308)	(.362, .379, .259)
11	uniform	uniform	(.411, .303, .286)
12	(.623, .377)	(.362, .333, .305)	uniform
13	uniform	uniform	(.367, .387, .246)
14	(.623, .377)	(.366, .333, .301)	(.428, .296, .276)
15	uniform	uniform	uniform
16	(.626, .374)	(.375, .333, .292)	(.379, .406, .215)
17	uniform	uniform	(.474, .279, .247)
18	(.638, .362)	(.395, .333, .272)	uniform
19	uniform	uniform	(.431, .487, .082)
20	(.750, .250)	(.500, .333, .167)	(.833, .139, .028)

Table 6: Mixing probabilities before symmetries are broken under Fixed Stopping. Bold numbers indicate probabilities that exceed uniform.

this stage, $v(19)$ is then computed as the expected utility of either action at the equilibrium value of α . The mixing probability in the previous stage can now be computed using this value and the maintained assumption of the rotation scheme. The equilibrium is determined recursively in this manner. Table 6 lists the mixing probabilities for each of the three games we study.

Notice first that there is an effect across rounds modulo the periodicity of the anticipated rotation or flipping scheme. In each game, equilibrium mixing is uniform when the number of rounds remaining is divisible by the periodicity of the continuation play. In other periods, mixing is non-uniform, and reflects the relative positions of the actions over the remaining finite horizon, conditional on symmetries being broken in that period. These positions are not necessarily ordinally the same as the stage game payoffs to the actions. Notice, e.g., the last column of Table 6 in which in the first round (and every three rounds thereafter) the middle action is played with higher probability than the high action.

7.8 Breaking symmetry and equilibrium mixing: Evidence

The first issue we address is whether subjects indeed play (totally) mixed strategies in the presence of symmetries. We do not observe mixed strategies directly, only their realizations, but there is strong evidence that subjects do indeed mix. Symmetry is in fact broken in almost all of the groups that play repeated allocation games. For instance, focusing on Fixed Stopping, there are but 2 groups out of 360 that never coordinate before the end of the match, receiving a total payoff of zero for the match. Under Random stopping, there

are more instances of groups never coordinating (24 out of 200 under Mixed Feedback and 51 out of 140 under Payoff Feedback), but most of these observations are simply due to the match having very few rounds. Unless actions are such that symmetry happens to be broken in the first stage of a match, the eventual breaking of symmetry, which typically occurs, requires at least some of the subjects to choose different pure actions across rounds in order for symmetry to eventually be broken.

For each treatment, under the presumption that subjects have in mind the rotation or flipping scheme that is empirically observed to be the dominant play in that session, our theory makes an explicit prediction for how subjects should mix before symmetry is broken. These are contained in Tables 5 and 6. In line with these predictions, we estimate a mixing distribution in the following way. For every session, we look at all rounds in which symmetry has not been broken, and we take the empirical frequencies over pure actions from these observations as an estimate of the mixing distribution. Aggregating the data in this way assumes stationarity across rounds, which is theoretically justified only under Random Stopping. We look for evidence of non-stationarity in the Fixed Stopping sessions below. This approach also assumes that subjects mix identically, as implied by symmetric equilibrium.

The data is summarized in Table 7. The panel labeled “All Rounds” lists the frequencies of action choices from all rounds in which symmetry had not yet been broken in the group. These distributions can be compared to the equilibrium mixed strategies computed above. For the Random Stopping sessions, this comparison is straightforward.

For the Fixed Stopping sessions, we observe first that there is no evidence in the data of the predicted periodicity effect. This can be seen by looking at the last panel of Table 7, labeled “Rounds with Uniform Prediction.” There is no systematic or significant difference between play in all rounds and play in those rounds in which the prediction is uniform.⁴³

For all sessions except session 1, Pearson chi-squared tests reject the hypothesis that the empirical frequencies come from the theoretical distribution of the model. But the differences are mostly qualitatively small. In most sessions, the main discrepancy is that H is overplayed relative to the prediction. This could be due to uncertainty in the continuation play that will be used. There is one notable exception: in session 2 L is overplayed. In fact, it is played more than half the time. This session has payoffs $x = (30, 1)$. One possible explanation, of the form studied first by Stahl and Wilson (1995) and taken up by Costa-Gomez, Crawford, and Broseta (2001) and Camerer, Ho, and Chong (2004), is that, anticipating a rotation scheme, and anticipating that an opponent is likely to play H (as is the case in other sessions), a level- k reasoner of the next level would optimally play L.

The estimated mixing distributions in Table 7 directly imply a probability of breaking symmetry in any given round. These figures are listed under “Empirical q ” in Table 8. These probabilities are maximized at the uniform distribution, and so they are bounded above by $\frac{1}{2}$ for 2-player games and by $\frac{6}{27}$ for 3-player games. Assuming independence across rounds, which empirically is not a bad approximation, any value of q induces a (geometric) distribution over the round in which symmetry is broken.

In the data we directly observe the breaking of symmetry. We use this data to obtain a value of q through maximum likelihood estimation. Some groups never break symmetry, and those observations are thus treated as truncated data. The likelihood function is easily derived as

$$\mathcal{L}(q) = -n \log(q) - \left(\sum_{m=1}^M n_m^* T_m + \sum_{g=1}^n (t_g - 1) \right) \log(1 - q),$$

⁴³The rounds included in the last panel of Table 7 depend on the session. In all 2-player games they are the rounds in which an even number of periods remain. In the 3-player games with $x = (30, 20, 10)$ they are again the rounds in which an even number of periods remain, since the flipping scheme has periodicity two. In all other 3-player games they are the rounds in which the number of remaining periods is divisible by three.

Session	Action Frequencies Before Symmetry is Broken							
	All Rounds				Rounds with Uniform Prediction			
	H	M	L	# Obs.	H	M	L	# Obs.
1	0.549		0.451	144				
2	0.486		0.514	208				
3	0.589		0.411	224	0.587		0.413	150
4	0.617		0.383	376	0.575		0.425	240
5	0.728		0.272	342	0.745		0.255	208
6	0.456	0.313	0.231	528				
7	0.402	0.363	0.235	498				
8	0.569	0.253	0.178	687				
9	0.375	0.406	0.219	753				
10	0.463	0.270	0.268	627				
11	0.329	0.362	0.309	741	0.292	0.380	0.328	411
12	0.407	0.422	0.171	666	0.408	0.440	0.152	375
13	0.483	0.262	0.255	726	0.497	0.246	0.256	195

Table 7: Aggregate action frequencies in rounds before symmetry is broken. “All Rounds” lists data from all such rounds. For sessions with Fixed Stopping, the right panel lists frequencies from the subset of rounds in which the theoretical prediction is uniform mixing.

where $t = (t_1, \dots, t_n)$ is a vector denoting the round at which symmetry was broken for every group that indeed broke symmetry at some point, T_m is the number of rounds that were played in match m , and n_m^* is the number of groups in match m that never broke symmetry. The results of this estimation are presented in the last panel of Table 8.

For the most part, the values of q estimated from the empirical mixing frequencies are very close to those estimated from the distribution over the round in which symmetry was broken.⁴⁴ Nor is there a systematic pattern of differences between the two estimation procedures. We take this as evidence that our estimation procedures are informative.

A striking feature of Table 8 is that, except in the Payoff Feedback sessions, the estimated values of q are very efficient. That is, they are close to $\frac{n!}{n^n}$, which is the value obtained under uniform mixing. This is precisely what is required for equilibrium mixing, given that we observe a preponderance of very nearly ex-post payoff-symmetric continuation play.

There is, however, variation across sessions with respect to how quickly subjects break symmetry. Our theory suggests that the sessions in which symmetry is broken most quickly will be those in which payoffs are most nearly symmetric. This is in fact precisely what we observe. For example, notice that the 3-player games with Payoff Feedback (sessions 6-8) have both low estimated values of q and high measures of payoff asymmetry, as subjects more frequently simply repeat the actions that led to positive payoffs. To express this relationship, we compare the measures of payoff symmetry S and S' from Table 2 with the estimated values

⁴⁴There is one session (session 1) in which groups broke symmetry at a rate that is statistically significantly faster than the theoretical upper bound of uniform mixing. We observe that with a relatively small number of subjects in the laboratory, there is one way to break symmetry faster. Having typically broken symmetry in their groups from the previous match, subjects can view themselves as entering the subsequent match (with a new randomly assigned partner) with a private type defined by the action she last played in her previous match. These types can map into first round actions in the obvious way. In a 2-player experiment with k subjects, even if all subjects perfectly used this convention, the probability of breaking symmetry in the first round is still bounded above by $\frac{k/2}{k-1} = 5/9$ for session 1 with $k = 10$. Thus, our best explanation is that subjects broke symmetry quickly partly by chance.

Session	Empirical q	Maximum Likelihood	
		q	$\mathcal{L}(q)$
1	0.495	0.653	46.49
2	0.500	0.452	71.61
3	0.484	0.536	77.35
4	0.473	0.426	128.22
5	0.396	0.404	115.32
6	0.198	0.165	78.76
7	0.206	0.199	82.79
8	0.154	0.118	83.06
9	0.200	0.171	114.95
10	0.200	0.187	100.58
11	0.221	0.198	123.04
12	0.176	0.225	118.42
13	0.193	0.206	123.28

Table 8: Coordination probabilities. “Empirical q ” lists probabilities induced by the distributions in Table 7. “Maximum Likelihood” lists estimates obtained from the observations of when symmetry is broken, with associated negative log likelihoods.

		Mixing efficiency	
		Empirical q	MLE q
Payoff	S	-0.237	-0.490
Symmetry	S'	-0.141	-0.413

Table 9: Correlations between payoff symmetry and mixing efficiency.

of q from Table 8. Table 9 shows the correlations between payoff symmetry and efficiency of mixing probabilities for all 3-player games (sessions 6-13). Our hypothesis is that these correlations should be negative. In all four cases, this prediction obtains in our data.⁴⁵ In summary, the evidence for how subjects play before symmetry is broken is very much in line with our theoretical predictions, based on Pareto efficiency and simplicity in the context of repeated allocation game.

8 Conclusion

We have sought to provide a theory of how to play certain repeated games. The theory elicits a focal point as a way to play the game suggested by the structure of the game itself. For such a theory to be meaningful, it should be based on principles that are broadly applicable, which we think of as conventions or meta-norms. The criteria that we leverage in this regard are Pareto optimality and simplicity. For the allocation games that we study, under the constraints of symmetric play, these criteria are capable of making an explicit prediction for equilibrium play in repeated allocation games.

This play consists of two phases. In the first phase, subjects mix symmetrically over pure actions until coordination is achieved. In the second phase, a continuation game ensues in

⁴⁵When looking at only the 3-player games with Mixed Feedback, the conclusion strengthens, with all correlations below -0.5 .

which play can be asymmetric, delivering continuation payoffs to the players conditional on the action they played when symmetry was broken.

As first suggested by Bhaskar (2000), our results highlight a general tradeoff between Pareto efficiency at the ex ante stage, and the extent of payoff asymmetry ex post. In particular, our finding that not all feasible payoffs can be supported by symmetric equilibria (Proposition 2) derives from the fact that ex-post payoff asymmetry induces inefficient mixing initially as players compete to obtain the higher continuation payoffs, resulting in costly delay. Thus, Pareto optimality dictates ex post symmetric payoffs (Corollary 1). In this sense, egalitarianism can be justified on the basis of efficiency alone, without reference to fairness per se.

The experiments that we run in the laboratory resoundingly support simplicity as a criterion to dictate how to deliver the efficient payoff profile.⁴⁶ Rotation schemes are, by far, the predominant type of continuation play. We explain this observation by virtue of the fact that rotation schemes deliver nearly symmetric payoffs, thereby incentivizing players to mix efficiently and minimize the delay cost imposed by symmetry.

Our analysis leaves open a number of questions. First, it would be interesting to explore the extent to which the link between payoff symmetry and Pareto efficiency extends to other symmetric games. The main intuition driving our results suggests that, in many contexts, there will be a tradeoff between payoff symmetry and ex ante efficiency. More generally, it would be useful to have a better understanding of the restrictions implicit in symmetric play in repeated games.

Second, it should be possible to extend our characterization of the Thue-Morse sequence, as the limit of sequences that award high payoffs to disadvantaged players, to games with more than two players. In particular, it would be interesting to know if its property of being most nearly symmetric for high discount factors is true more generally. Also, Proposition 6 could potentially be strengthened so as to dispense with the requirement of periodicity. Finally, it is potentially interesting to understand if there are other repeated games for which Pareto efficiency at high discount factors implies similar play to that described by the Thue-Morse sequence.

A Proof of Proposition 2

Fix an allocation game Γ . Let $X = \sum_i x_{a_i}$. Fix any payoff profile $z \in \mathcal{P}_\Gamma$ on the Pareto frontier of Γ that is not completely symmetric, i.e., for which there exist i and j with $z_i \neq z_j$. In fact, let, without loss of generality, $z_1 \geq \dots \geq z_n$. Fix an $\epsilon > 0$ and, finally, consider any payoff profile $w \in \mathcal{U}_\epsilon^z$.

We will demonstrate that ϵ can be chosen small enough such that for every δ large enough to allow $w \in \mathcal{F}^z(\delta)$, and every symmetric Markov strategy profile σ that delivers payoffs w , it must be that σ is not an equilibrium.

Since σ is Markovian it induces a particular probability of coordination, q , at every stage in which symmetries are unbroken. It prescribes, as well, a given \bar{w} as the continuation payoff profile upon breaking of symmetries. The most stringent case to consider is when σ prescribes an efficient (always coordinated) continuation play upon the breaking of symmetries, so that $\sum_i \bar{w}_i = X$. Defining

$$V(q) = q \sum_t (1-q)^t \delta^t = \frac{q}{1-\delta(1-q)},$$

we have, by definition, that $w = \bar{w}V(q)$.

⁴⁶The experiments in fact support simplicity as a criterion that comes after efficiency. Players do not use the simplest strategies. They instead use the simplest strategies among those that are close to efficient.

Consider a unilateral deviation from σ to the strategy σ' that plays a_1 with probability one at every history before symmetries are broken, and then plays according to σ in the continuation. For a contradiction, we want to show ϵ can be chosen small enough such that this deviation is profitable for every δ at which w is feasible.

Observe that $q = n!(1 - \sum_{i>1} \sigma_{a_i}) \prod_{i>1} \sigma_{a_i}$. Similarly, the coordination probability induced when one player, say i , uses σ' and the other players all use σ , is $q' = (n-1)! \prod_{i>1} \sigma_{a_i}$. Thus, $q = n(1 - \sum_{i>1} \sigma_{a_i})q'$ implying, since σ is totally mixed, that $q' > q/n$.

For $w \in \mathcal{U}_\epsilon^z$ it must be that $V(q) \geq 1 - \epsilon$, i.e.,

$$q \geq \frac{(1 - \delta)(1 - \epsilon)}{1 - \delta(1 - \epsilon)} \equiv T.$$

The intuition is that the deviation to σ' is profitable because feasibility of w requires δ to be large enough that the reduction in coordination probability from q to q' is not too costly, whereas the gain, upon the breaking of symmetries, is significant.

We have

$$\begin{aligned} (1) \quad u(\sigma'_i, \sigma_{-i}) &= \bar{w}_{a_1} V(q') > \bar{w}_{a_1} V(q/n) \geq \bar{w}_{a_1} V(T/n) \\ (2) \quad &> \bar{w}_{a_1} \frac{1}{n(1 - \delta) + \delta} \\ (3) \quad &> \left(\frac{1}{n} \sum_i \bar{w}_{a_i}\right) V(q) \\ (4) \quad &= u(\sigma). \end{aligned}$$

(1) holds because V is increasing. (2) is true for sufficiently small ϵ because $\lim_{\epsilon \rightarrow 0} V(T/n) = \frac{1}{n(1 - \delta) + \delta}$. There exists a $\bar{\delta}$ so that (3) holds for all $\delta > \bar{\delta}$, since as $\lim_{\delta \rightarrow 1} \frac{1}{n(1 - \delta) + \delta} = 1 \geq V(q)$. Since $\bar{w}_{a_1} > \frac{1}{n} \sum_i \bar{w}_{a_i}$ by a given constant amount, ϵ and $\bar{\delta}$ can be chosen without reference to q , achieving the desired contradiction. Finally, notice that as $\epsilon \rightarrow 0$, feasibility requires that δ converge to one, while $\bar{\delta}$ can remain bounded away from one, completing the proof.

B Proof of Proposition 3

Let $\bar{\mathcal{E}}(\bar{w}) \subset \mathcal{F}$ be such that any $(w^H, w^L) \in \bar{\mathcal{E}}(\bar{w})$ satisfies $w^H, w^L \leq \bar{w}$. We have $\bar{\mathcal{E}}(1) = \mathcal{F}$, but for $\bar{w} < 1$ $\bar{\mathcal{E}}(\bar{w})$ is a proper subset of \mathcal{F} .

The proof proceeds by showing that whenever $\bar{w} > \frac{3}{4}$, there is a $\bar{w}' < \bar{w}$ such that $f(\bar{\mathcal{E}}(\bar{w})) \subset \bar{\mathcal{E}}(\bar{w}')$. This directly implies that any $G \in \mathcal{G}$ with a payoff greater than $\frac{3}{4}$ cannot be a fixed point of f .

Let $(w_{HH}^H, w_{HH}^L) \in \bar{\mathcal{E}}(\bar{w})$ be the continuation after HH and $(w_{LL}^H, w_{LL}^L) \in \bar{\mathcal{E}}(\bar{w})$ be the continuation after LL . Furthermore let $(w^H, w^L) \in \mathcal{F}$ be the (immediately paid out) continuation after HL and LH . Let $\alpha \in [0, 1]$ denote the probability players attach to pure action H in stage 0, which must be the same for both players by symmetry.

At stage 0, when players contemplate their choice of α , expected payoffs from choosing pure action H and L are given by

$$u(H, \alpha) = \frac{w_{HH}^H + w_{HH}^L}{2} \delta \alpha + (1 - \alpha) w^H,$$

and

$$u(L, \alpha) = w^L \alpha + (1 - \alpha) \frac{w_{LL}^H + w_{LL}^L}{2} \delta.$$

This follows from the continuation payoffs and the fact that after HH and LL both players are equally likely, given attainability, to end up being the H -player or L -player at the moment when symmetries are broken.

Players now choose α such that neither of them has an incentive to deviate to another (mixed) action. Thus, we are looking for a symmetric equilibrium of the following symmetric 2×2 game.

	H	L
H	$\delta \frac{w_{HH}^H + w_{LL}^L}{2}, \delta \frac{w_{HH}^H + w_{LL}^L}{2}$	w^H, w^L
L	w^L, w^H	$\delta \frac{w_{LL}^H + w_{LL}^L}{2}, \delta \frac{w_{LL}^H + w_{LL}^L}{2}$

Note that if this game has only pure equilibria (such as H being a dominant strategy) then the ex-ante, at time 0, expected payoff must be less than or equal to $\delta \bar{w} < \bar{w}$. Thus the desired conclusion would hold in fact for any $\bar{w} > 0$.

So, the interesting case involves continuations such that this game has a strictly mixed equilibrium.

The unique completely mixed symmetric equilibrium is given by

$$\alpha^* = \frac{2w^H - (w_{LL}^H + w_{LL}^L)\delta}{2(w^H + w^L) - (w_{HH}^H + w_{HH}^L + w_{LL}^H + w_{LL}^L)\delta}.$$

Given the continuation profile and the induced α^* we then have that the ex-ante, at stage 0, expected payoff to the (eventual) H -player at the (possibly later) moment symmetries are broken is given by

$$(5) \quad w^* = (\alpha^*)^2 w_{HH}^H \delta + 2\alpha^*(1 - \alpha^*)w^H + (1 - \alpha^*)^2 w_{LL}^H \delta.$$

We are now trying to show that $w^* < \bar{w}$ whenever $\bar{w} > \frac{3}{4}$, given $w_{HH}^H, w_{LL}^H \leq \bar{w}$, as well as the aforementioned restrictions on the continuation profile, and subject to the incentive constraints.

In order to do so we distinguish two cases. Suppose first that $\alpha^* \leq \frac{1}{2}$. That is,

$$\frac{2w^H - (w_{LL}^H + w_{LL}^L)\delta}{2(w^H + w^L) - (w_{HH}^H + w_{HH}^L + w_{LL}^H + w_{LL}^L)\delta} \leq \frac{1}{2},$$

or, equivalently,

$$2w^H - (w_{LL}^H + w_{LL}^L)\delta \leq 2w^L - (w_{HH}^H + w_{HH}^L)\delta.$$

Thus,

$$w^H - w^L \leq \frac{1}{2} [(w_{LL}^H + w_{LL}^L)\delta - (w_{HH}^H + w_{HH}^L)\delta].$$

Given $w^H + w^L \leq 1$, we finally have,

$$w^H \leq \frac{1}{2} + \frac{1}{4}\delta \leq \frac{3}{4}.$$

From Equation 5, as long as $w^H \leq \bar{w}\delta$ we must have $w^* \leq \bar{w}\delta$ as well. This is definitely true if $\frac{1}{2} + \frac{1}{4}\delta \leq \bar{w}\delta$, i.e. if $\bar{w} \geq \frac{\frac{1}{2} + \frac{1}{4}\delta}{\delta}$. Thus, if $\bar{w} > \frac{3}{4}$ (in the limit when $\delta \rightarrow 1$), we have $w^* < \bar{w}$, as desired.

Suppose now that $\alpha^* > \frac{1}{2}$. Obviously,

$$w^* \leq w^{**} = \max_{\alpha \in [\frac{1}{2}, 1], w^H, w_{HH}^H, w_{LL}^H} \alpha^2 w_{HH}^H \delta + 2\alpha(1 - \alpha)w^H + (1 - \alpha)^2 w_{LL}^H \delta,$$

subject to the given restrictions on w^H , w_{HH}^H , and w_{LL}^H . Further, the indifference condition⁴⁷ required to induce mixing implies that $w^H \leq 1 - \frac{w_{HH}^H + w_{HH}^L}{2}\delta$. w^{**} is increasing in each of

⁴⁷There are two possibilities for a mixed equilibrium: either the game is a coordination game or of the Hawk-Dove variety. For a coordination game we must have $w^H \leq \frac{w_{LL}^H + w_{LL}^L}{2}\delta < \frac{1}{2}$ and $w^L \leq \frac{w_{HH}^H + w_{HH}^L}{2}\delta < \frac{1}{2}$. For a game to be of the Hawk-Dove variety we must have $w^H \geq \frac{w_{LL}^H + w_{LL}^L}{2}\delta$ and $w^L \geq \frac{w_{HH}^H + w_{HH}^L}{2}\delta$, which, by $w^H + w^L \leq 1$, implies that $w^H \leq 1 - \frac{w_{HH}^H + w_{HH}^L}{2}\delta > \frac{1}{2}$ (as long as δ sufficiently close to 1). Thus, for w^H the Hawk-Dove case is less restrictive than the coordination case.

w^H, w_{HH}^H, w_{LL}^H . Thus, $w^{**} \leq \max_{\alpha \in [\frac{1}{2}, 1], w_{HH}^H \in [0, \bar{w}]} \alpha^2 w_{HH}^H \delta + 2\alpha(1 - \alpha) \left(1 - \frac{w_{HH}^H}{2} \delta\right) + (1 - \alpha)^2 \bar{w} \delta$. But for $\alpha > \frac{1}{2}$, this expression is maximized at $w_{HH}^H = \bar{w}$. Thus we have

$$w^{**} \leq \max_{\alpha \in [\frac{1}{2}, 1]} \alpha^2 \bar{w} \delta + 2\alpha(1 - \alpha) \left(1 - \frac{\bar{w}}{2} \delta\right) + (1 - \alpha)^2 \bar{w} \delta.$$

Therefore, $w^{**} \leq \bar{w} \delta$ if $\left(1 - \frac{\bar{w}}{2} \delta\right) < \bar{w} \delta$. This, in turn, is true if $\bar{w} \delta > \frac{2}{3}$.

To summarize, for δ close enough to 1, $\bar{w} > \frac{3}{4}$ ensures that in order for a set $\mathcal{E}^a(\delta)$ to be a fix point of the mapping f it has to satisfy $\mathcal{E}^a(\delta) \subset \bar{\mathcal{E}}\left(\frac{3}{4}\right)$. QED

C Proof of Proposition 4

Note that each stationary symmetric strategy profile has associated with it a (normalized discounted) payoff \bar{w}^H , the continuation payoff to the player who plays H when symmetries are broken. \bar{w}^H must be consistent with the initial mixing probability α . Notice that, provided $\delta \geq \frac{1}{2}$, any continuation $\bar{w}^H \in [0, 1]$ is feasible.⁴⁸ Given α , at any period in which players are symmetric, the probability of the symmetry breaking is $q(\alpha) = 2\alpha(1 - \alpha)$. Ex ante, when mixing initially with α and playing an efficient continuation once symmetries are broken, each player has expected payoff

$$u(\alpha, \delta) = \frac{1}{2} \sum_{t=0}^{\infty} q(1 - q)^t \delta^t = \frac{q}{2(1 - \delta(1 - q))}.$$

In order to incentivize the players to mix initially, we must have equal expected payoffs from either action, which requires

$$\alpha \delta u + (1 - \alpha) \bar{w}^H = (1 - \alpha) \delta u + \alpha(1 - \bar{w}^H).$$

On the left hand side, when choosing H , there are two possibilities. With probability α the other player chooses H , symmetries are not broken, and the game continues at the next date with continuation payoff u . On the other hand, with probability $(1 - \alpha)$ the other player chooses L , in which case symmetry is broken and the continuation \bar{w}^H is realized. Similarly, when choosing L , with probability $(1 - \alpha)$ symmetry is not broken, and when it is, a continuation of $1 - \bar{w}^H$ is realized. Solving, we obtain

$$\bar{w}^H(\alpha, \delta) = \frac{\alpha(1 - \delta\alpha)}{1 - \delta(1 - q(\alpha))}.$$

Even though players both expect u at the beginning of the game, ex post they may obtain different payoffs. We will show which payoff profiles (w^H, w^L) can be supported by symmetric Markov equilibria.

We have that

$$w^H(\alpha, \delta) = \bar{w}^H \sum_{t=0}^{\infty} q(1 - q)^t \delta^t = \bar{w}^H \frac{q}{1 - \delta(1 - q)}$$

and

$$w^L(\alpha, \delta) = (1 - \bar{w}^H) \sum_{t=0}^{\infty} q(1 - q)^t \delta^t = (1 - \bar{w}^H) \frac{q}{1 - \delta(1 - q)}.$$

For a given δ , these expressions trace out a parametric curve of equilibrium payoff profiles as α varies from zero to one, starting and ending at the origin. Notice that, given a pair

⁴⁸This follows from Sorin (1986), and appears as (Mailath and Samuelson 2007, Lemma 7.3.1).

(w^H, w^L) , any payoff along the ray pointing to the origin is also supportable, by using an inefficient continuation that gives total payoff $F < 1$. We would then have $u = \frac{F}{2} \frac{q}{1-\delta(1-q)}$ and continuation payoffs once symmetries are broken of $(\bar{w}^H, F - \bar{w}^H)$. Thus, the region defined by the parametric curve represents the set of symmetric Markov equilibrium payoffs.

We now show that as δ tends to one, one third of feasible payoffs are supported by a symmetric Markov equilibria. First notice that $\lim_{\delta \rightarrow 1} (w^H(\frac{1}{2}, \delta), w^L(\frac{1}{2}, \delta)) = (\frac{1}{2}, \frac{1}{2})$. Define next the area

$$A' = \lim_{\delta \rightarrow 1} \int_{x=0}^{\frac{1}{2}} w^H(\alpha, \delta) \frac{\partial w^L(\alpha, \delta)}{\partial \alpha} d\alpha,$$

which gives the limiting area under the upper lobe of equilibrium payoffs. Thus, the limiting area sustained by symmetric Markov equilibrium payoffs is

$$A = 2(A' - \frac{1}{8}) = \frac{1}{6},$$

where the area A' is computed with some straightforward but tedious calculus and algebra. QED

D An example of a Pareto optimal equilibrium that fails strong ex post symmetry

In this appendix we study a particular symmetric repeated game that falls outside the class of allocation games. Nevertheless it can be analyzed in much the same way as an allocation game, demonstrating our first point of this section that much of our analysis extends somewhat outside the games we study above. Second, this example conveys the fact that the symmetry structure in a repeated game can be more subtle than that pertaining to allocation games.

The example is a 3-player game in which all players are symmetric. However they are symmetric in a specific way. If, e.g., we observe one player choosing one action, say H , and the other two both choosing another action, say L then **all** symmetries are nevertheless immediately broken. This was not the case for 3-player allocation games. In such a game the action profile (H, L, L) would not have resulted in symmetry breaking. Even if monitoring was perfect in the allocation game and the action profile was common knowledge, the history (H, L, L) would allow us to differentiate only between the H -player and the two L -players, but not between the two L -players. In the example we present here, the two L -players are further differentiated by their “relative position” to the H -player, as will be explained below. The third point of this appendix is to show that there are games, such as the one given here, in which efficient symmetric equilibria of the repeated game, while they have to be ex-post payoff-symmetric, do not have to be strongly ex-post payoff-symmetric.

Consider the following 3-person stage game, where $1 < x < 3$.

	H			M			L		
	<i>H</i>	<i>M</i>	<i>L</i>	<i>H</i>	<i>M</i>	<i>L</i>	<i>H</i>	<i>M</i>	<i>L</i>
<i>H</i>	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	3,1,x	0,0,0	0,0,0	0,0,0
<i>M</i>	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	x,3,1	0,0,0	0,0,0
<i>L</i>	0,0,0	1,x,3	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0	0,0,0

This game can be thought of as another natural extension of the 2-player Battle-of-the-Sexes game. In this stage game there are three symmetries⁴⁹ in the sense of Nash (1951)

⁴⁹Loosely speaking, a symmetry of a game, here for players, is simply a set of permutations of player names that leaves the matrix depiction of the game unchanged. One such permutation is, of course, the identity mapping.

(see also Alos-Ferrer and Kuzmics (2008)) of this game: $\{(1 \rightarrow 1, 2 \rightarrow 2, 3 \rightarrow 3), (1 \rightarrow 2, 2 \rightarrow 3, 3 \rightarrow 1), (1 \rightarrow 3, 3 \rightarrow 2, 2 \rightarrow 1)\}$.

Now consider the repeated game with this stage game, in which monitoring is perfect. Given this and the presence of only three symmetries (instead of the possible 6), any history that identifies one of the players necessarily identifies all three players, breaking all symmetries at once. To see that, for instance, (realized) action profile (H, L, L) breaks all symmetries, note first that the H -player is of course identified. The two L -players can be differentiated in terms of what action they would need to play in order for an action profile in which the H -player plays H to result in positive payoffs to all. Of all action profiles in which, for instance, player 1 chooses H only (H, L, M) results in positive profits. In particular (H, M, L) results in zero payoffs to all players. Thus, the two L -players can be distinguished (one could call them the “prospective” L and M -players).

The only histories, then, after which symmetries remain are (H, H, H) , (M, M, M) , and (L, L, L) .

Now consider the following convention, which specifies actions in any continuation game. After the first occurrence of any permutation of (L, M, H) , continue with something that is ex post payoff-symmetric (and on the Pareto frontier) by using time-averaging. After a stage in which only one of the actions is chosen by exactly one player, coordinate forever after on that player’s preferred outcome. After any other profile (i.e. those for which symmetries remain), continue mixing as in the previous period. It remains to be shown what mixing probabilities are induced by this continuation play at a symmetric equilibrium.

Notice that this convention is not strongly ex-post payoff-symmetric, as it allows for the positive probability outcome of awarding per-period continuation payoffs of (any permutation of) $(1, x, 3)$.

The mixing probabilities that minimize the expected time until symmetries are broken are those that minimize the probability of realizing one of the profiles (H, H, H) , (M, M, M) , and (L, L, L) . Thus, the efficient mixing probabilities are $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

Our claim is that the initial equilibrium mixing implied by this protocol is also $\alpha = (\alpha_H, \alpha_M, \alpha_L) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. For this to be true, it must be that for every action, conditional on symmetries being broken, the expected continuation values are identical. That is to say, the strategy profile needs to be ex-post payoff-symmetric.

The argument for the three actions are analogous, so without loss of generality consider player 1 taking action L , and the other players choose actions according to α . Conditional on symmetry being broken, there are essentially four possibilities to consider. First, it could be that the other players play (M, M) or (H, H) , in which case player 1’s continuation payoff is 3. This happens with probability $\frac{2}{9}$. Second, it could be the case that the other players play (M, H) or (H, M) , in which case player 1’s continuation payoff is $(1 + x + 3)/3$. This happens with probability $\frac{2}{9}$ as well. Third, it could be the case that player 2 plays L and player 3 plays M or H , in which case player 1’s continuation payoff is 1. This happens with probability $\frac{2}{9}$. Finally, player 3 could play L and player 2 could play M or H , in which case player 1’s continuation payoff is x . This also happens with probability $\frac{2}{9}$.

In summary, for any action played, there is a $\frac{1}{9}$ chance that symmetry is not broken. If, however, they are broken, there are four equally likely possibilities that involve continuation payoffs of 1, x , 3, or $(1 + x + 3)/3$. Thus, uniform mixing is played in equilibrium, as claimed.

The ability to construct such a protocol requires that we be able to treat the three actions in a symmetric way in the continuation games after symmetries are broken. If we are to construct a protocol that is not strongly ex-post payoff-symmetric, then to achieve Pareto-optimal mixing, it must be that these asymmetries are constructed in a balanced way in order to equalize incentives for the initial mixing. In some settings this may not be possible, in which case the only Pareto-optimal equilibria involve strong ex -ost payoff-symmetric outcomes.

References

- ALLOUCHE, J.-P., AND J. SHALLIT (1999): “The Ubiquitous Prouhet-Thue-Morse Sequence,” in *Sequences and Their applications, Proc. SETA '98*, ed. by T. Ding, C. Hellesest, and H. Niederreiter, pp. 1–16. Springer-Verlag, New York.
- ALOS-FERRER, C., AND C. KUZMICS (2008): “Hidden symmetries and focal points,” Mimeo.
- AUMANN, R. J. (1981): “Survey of repeated games,” in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, pp. 11–42. Bibliographisches Institut Mannheim / Wien / Zürich.
- BARON, D., AND E. KALAI (1993): “The Simplest Equilibria of a Majority Division Game,” *Journal of Economic Theory*, 61, 290–301.
- BEN PORATH, E. (1986): “Repeated games with bounded complexity,” Mimeo, Stanford University.
- BHASKAR, V. (2000): “Egalitarianism and Efficiency in Repeated Symmetric Games,” *Games and Economic Behavior*, 32, 247–262.
- BLUME, A. (2000): “Coordination and Learning With a Partial Language,” *Journal of Economic Theory*, 95, 1–36.
- BLUME, A., AND U. GNEEZY (2000): “An Experimental Investigation of Optimal Learning in Coordination Games,” *Journal of Economic Theory*, 90, 161–172.
- (2010): “Cognitive Forward Induction and Coordination Without Common Knowledge: An Experimental Study,” *Games and Economic Behavior*, 68, 488–511.
- CAMERER, C., T. HO, AND J. CHONG (2004): “A Cognitive Hierarchy Model of Behavior in Games,” *Quarterly Journal of Economics*, 119(3), 861–898.
- CASAJUS, A. (2000): “Focal Points in Framed Strategic Forms,” *Games and Economic Behavior*, 32, 263–291.
- (2001): *Focal Points in Framed Games*. Springer.
- CASON, T., S.-H. P. LAU, AND V.-L. MUI (2012): “Learning, Teaching, and Turn Taking in the Repeated Assignment Game*,” Mimeo.
- CHATTERJEE, K., AND H. SABOURIAN (2009): “Game Theory and Strategic Complexity,” in *Encyclopedia of Complexity and System Science*, ed. by R. A. Meyers. Springer.
- COSTA-GOMEZ, M., V. CRAWFORD, AND B. BROSETA (2001): “Cognition and Behavior in Normal-form Games: An experimental study,” *Econometrica*, 69(5), 1193–1235.
- CRAWFORD, V., AND H. HALLER (1990): “Learning How to Cooperate: Optimal Play in Repeated Coordination Games,” *Econometrica*, 58, 571–596.
- FARRELL, J. (1987): “Cheap Talk, Coordination and Entry,” *Rand Journal of Economics*, 18, 34–39.
- FUDENBERG, D., AND E. MASKIN (1991): “On the Dispensability of Public Randomizations in Discounted Repeated Games,” *Journal of Economic Theory*, 53, 428–438.
- HARSANYI, J., AND R. SELTEN (1988): *General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- KALAI, E. (1990): “Bounded Rationality and Strategic Complexity in Repeated Games,” in *Game Theory and Applications*, ed. by T. Ichiishi, A. Neyman, and Y. Tauman, pp. 131–157. San Diego: Academic Press.
- KALAI, E., AND W. STANFORD (1988): “Finite Rationality and Interpersonal Complexity in Repeated Games,” *Econometrica*, 56 (2), 387–410.

- KUZMICS, C., AND B. ROGERS (2012): “A Comment on V. Bhaskar (2000) ”Egalitarianism and Efficiency in Repeated Symmetric Games”, *Games and Economic Behavior*, 32: 247-262,” *Games and Economic Behavior*, 74(1), 240–242.
- LAU, S.-H. P., AND V.-L. MUI (2008): “Using Turn Taking to Mitigate Coordination and Conflict Problems in the Repeated Battle of the Sexes Game,” *Theory and Decision*, 65, 153–183.
- (2012): “Using Turn Taking to Achieve Intertemporal Cooperation and Symmetry in Infinitely Repeated 2x2 Games,” *Theory and Decision*, 72, 167–188.
- LUCE, R., AND H. RAIFFA (1957): *Games and Decisions: Introduction and Critical Survey*. Wiley, New York, NY.
- MAILATH, G., AND L. SAMUELSON (2007): *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press, Oxford, UK.
- MORSE, M. (1921): “Recurrent Geodesics on a Surface of Negative Curvature,” *Trans. Amer. Math. Soc.*, 22, 84–100.
- NASH, J. (1951): “Non-Cooperative Games,” *Annals of Mathematics*, 52, 286–295.
- NEYMAN, A. (1985): “Bounded complexity justifies cooperation in the finitely-repeated Prisoners’ Dilemma,” *Economics Letters*, 19, 227–229.
- PRISBEY, J. (1992): “An Experimental Analysis of Two-Person Reciprocity Games,” Mimeo, Caltech Working Paper 787.
- RUBINSTEIN, A. (1986): “Finite automata play the repeated Prisoners’ Dilemma,” *Journal of Economic Theory*, 39 (1), 83–96.
- SCHELLING, T. C. (1960): *The Strategy of Conflict*. Harvard University Press, Cambridge.
- SORIN, S. (1986): “On repeated games with complete information,” *Mathematics of Operations Research*, 11, 147–160.
- STAHL, D., AND P. WILSON (1995): “On Players’ Models of Other Players: Theory and experimental evidence,” *Games and Economic Behavior*, 10, 218–254.
- THUE, A. (1906): “Über unendliche Zeichenreihen,” *Norske vid. Selsk. Skr. Mat. Nat. Kl.*, 7, 1–22, Reprinted in Selected Mathematical Papers of Axel Thue (Ed. T. Nagell). Oslo: Universitetsforlaget, pp. 139-158, 1977.
- (1912): “Über die gegenseitige Lage gleicher Teile gewisser Zeichenreihen,” *Norske vid. Selsk. Skr. Mat. Nat. Kl.*, 1, 1–67, Reprinted in Selected Mathematical Papers of Axel Thue (Ed. T. Nagell). Oslo: Universitetsforlaget, pp. 413-478, 1977.